



Fabric OS

Administrator's Guide

Supporting Fabric OS v5.0.1

**Supporting SilkWorm 200E, 3014, 3016, 3250, 3850, 3900, 4012,
4100, 12000, 24000, 48000**


Publication Number: 53-0000518-09

Publication Date: 07/11/05

Copyright © 2005, Brocade Communications Systems, Incorporated.

ALL RIGHTS RESERVED.

Publication Number: 53-0000518-09

Brocade, the Brocade B weave logo, Secure Fabric OS, and SilkWorm are registered trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. FICON, IBM  BladeCenter are registered trademarks of IBM Corporation in the U.S. and other countries. All other brands, products, or service names are or may be trademarks or service marks of, and are used to identify, products or services of their respective owners.

Notice: The information in this document is provided “AS IS,” without warranty of any kind, including, without limitation, any implied warranty of merchantability, noninfringement or fitness for a particular purpose. Disclosure of information in this material in no way grants a recipient any rights under Brocade's patents, copyrights, trade secrets or other intellectual property rights. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use.

The authors and Brocade Communications Systems, Inc. shall have no liability or responsibility to any person or entity with respect to any loss, cost, liability, or damages arising from the information contained in this book or the computer programs that accompany it.

Notice: The product described by this document may contain “open source” software covered by the GNU General Public License or other open source license agreements. To find-out which open source software is included in Brocade products, view the licensing terms applicable to the open source software, and obtain a copy of the programming source code, please visit <http://www.brocade.com/support/oscd>.

Export of technical data contained in this document may require an export license from the United States Government.

Brocade Communications Systems, Incorporated

Corporate Headquarters

Brocade Communications Systems, Inc.
1745 Technology Drive
San Jose, CA 95110
Tel: 1-408-333-8000
Fax: 1-408-333-8101
E-mail: info@brocade.com

Asia-Pacific Headquarters

Brocade Communications Singapore Pte. Ltd.
9 Raffles Place
#59-02 Republic Plaza 1
Singapore 048619
Tel: +65-6538-4700
Fax: +65-6538-0302
Email: apac-info@brocade.com

European and Latin American Headquarters

Brocade Communications Switzerland Sàrl
Centre Swissair
Tour A - 2ème étage
29, Route de l'Aéroport
Case Postale 105
CH-1215 Genève 15
Switzerland
Tel: +41 22 799 56 40
Fax: +41 22 799 56 41
Email: emea-info@brocade.com

Document History

The following table lists all versions of the *Fabric OS Administrator's Guide*.

Document Title	Publication Number	Summary of Changes	Publication Date
<i>Fabric OS Procedures Guide</i>	53-0000518-02	First released edition.	April 2003
<i>Fabric OS Procedures Guide</i>	53-0000518-03	Revised for Fabric OS v4.2.0.	December 2003
<i>Fabric OS Procedures Guide</i>	53-0000518-04	Revised to include switch-specific information.	March 2004
<i>Fabric OS Procedures Guide</i>	53-0000518-05	Revised for Fabric OS v4.4.0.	September 2004
<i>Fabric OS Procedures Guide</i>	53-0000518-06	Revised to add RADIUS and SSL procedures.	October 2004
<i>Fabric OS Administrator's Guide</i>	53-0000518-07	Revised book title. Added information about 200E, 4012, and 48000 switches.	April 2005
<i>Fabric OS Administrator's Guide</i>	53-0000518-09	Added release note updates.	July 2005

Contents

About This Document

How This Document Is Organized	xv
Supported Hardware and Software	xvi
What's New in This Document.	xvii
Document Conventions.	xvii
Additional Information	xviii
Getting Technical Help	xx
Document Feedback	xxi

Section I **Standard Features**

Chapter 1 **Introducing Fabric OS CLI Procedures**

About Procedural Differences.	1-1
Scope and References	1-2
About the CLI.	1-2
Help Information.	1-3
Displaying Command Help	1-3
Displaying Additional Help Topics	1-4

Chapter 2 **Performing Basic Configuration Tasks**

Connecting to the Command Line Interface.	2-1
Setting the IP Address.	2-3
Setting the Default Account Passwords	2-3
Setting the Date and Time.	2-5
Maintaining Licensed Features.	2-8
Customizing the Switch Name	2-10
Customizing the Chassis Name	2-12

Disabling and Enabling a Switch	2-12
Disabling and Enabling a Port	2-12
Activating Ports on Demand.....	2-13
Making Basic Connections	2-14
Connecting to Devices	2-15
Connecting to Other Switches	2-15
Working With Domain IDs	2-15
Linking Through a Gateway	2-17
Checking Status	2-18
Tracking and Controlling Switch Changes	2-20

Chapter 3 Configuring Standard Security Features

Secure Protocols	3-1
Ensuring Network Security.....	3-3
Configuring the Telnet Interface.....	3-3
Blocking Listeners	3-5
Accessing Switches and Fabrics.....	3-5
Creating and Maintaining User-Defined Accounts.....	3-6
Changing an Account Password.....	3-8
Setting Up RADIUS AAA Service.....	3-9
Configuring the RADIUS Server	3-10
Configuring the Switch.....	3-15
Enabling and Disabling Local Authentication	3-18

Configuring for the SSL Protocol	3-18
Browser and Java Support	3-18
Summary of SSL Procedures	3-19
Choosing a Certificate Authority	3-20
Generating a Public/Private Key	3-20
Generating and Storing a CSR	3-20
Obtaining Certificates	3-21
Installing a Switch Certificate	3-22
Activating a Switch Certificate	3-22
Configuring the Browser	3-23
Installing a Root Certificate to the Java Plug-in	3-24
Displaying and Deleting Certificates	3-24
Troubleshooting Certificates	3-25
Configuring for SNMP	3-25
Setting the Security Level	3-26
Using the snmpConfig Command	3-27
Using Legacy Commands for SNMPv1	3-30
Configuring Secure File Copy	3-34
Setting the Boot PROM Password	3-34
With a Recovery String	3-35
Without a Recovery String	3-37
Recovering Forgotten Passwords	3-39

Chapter 4 **Maintaining Configurations and Firmware**

Maintaining Configurations	4-1
Displaying Configuration Settings	4-1
Backing Up a Configuration	4-1
Restoring a Configuration	4-2
Restoring Configurations in a FICON Environment	4-4
Downloading Configurations Across a Fabric	4-5
Printing Hard Copies of Switch Information	4-5

Maintaining Firmware	4-5
Obtaining and Unzipping Firmware	4-6
Checking Connected Switches	4-6
About the Download Process	4-7
Upgrading SilkWorm Switches	4-8
Upgrading SilkWorm Directors	4-11
Troubleshooting Firmware Downloads	4-15

Chapter 5 Configuring SilkWorm 12000, 24000, and 48000 Directors

Identifying Ports	5-1
By Slot and Port Number	5-2
By Port Area ID	5-2
Basic Blade Management	5-3
Powering Port Blades Off and On	5-3
Disabling and Enabling Port Blades	5-3
Conserving Power	5-4
Blade Terminology and Compatibility	5-4
CP Blades	5-6
Port Blade Compatibility	5-6
Setting Chassis Configurations	5-6
Obtaining Slot Information	5-8
Configuring a New SilkWorm 24000 with Two Domains	5-9
Converting an Installed SilkWorm 24000 to Support Two Domains	5-10
Combining SilkWorm 12000 and 24000 Port Blades in One Chassis ...	5-11
Setting the Blade Beacon Mode	5-12

Chapter 6 Routing Traffic

About Data Routing and Routing Policies	6-1
Specifying the Routing Policy	6-2
Assigning a Static Route	6-3
Specifying Frame Order Delivery	6-3

Using Dynamic Load Sharing.	6-4
Viewing Routing Path Information.	6-5
Viewing Routing Information Along a Path.	6-7

Chapter 7 Administering FICON Fabrics

FICON Overview	7-1
Configuring Switches	7-4
Preparing a Switch	7-4
Configuring a Single Switch.	7-5
Configuring a High-Integrity Fabric.	7-5
Setting a Unique Domain ID	7-6
Displaying Information.	7-7
Link Incidents	7-7
Registered Listeners	7-7
Node Identification Data.	7-7
FRU Failures.	7-7
Swapping Ports	7-8
Clearing the FICON Management Database	7-8
Using FICON CUP.	7-9
Setup Summary.	7-9
Enabling and Disabling FICON Management Server Mode	7-9
Displaying the fmsmode Setting.	7-11
Displaying Mode Register Bit Settings	7-12
Setting Mode Register Bits.	7-13
Persistently Enabling/Disabling Ports	7-13
Port and Switch Naming Standards	7-14
Adding and Removing FICON CUP Licenses	7-14
Zoning and PDCM Considerations.	7-15
Backing Up and Restoring Configurations.	7-15
Troubleshooting	7-15
Identifying Ports	7-16
Backing Up FICON Files	7-16

	Recording Configuration Information	7-17
Chapter 8	Configuring the Distributed Management Server	
	Enabling and Disabling the Platform Services	8-1
	Controlling Access	8-2
	Configuring the Server Database	8-5
	Controlling Topology Discovery	8-6
Chapter 9	Working With Diagnostic Features	
	Viewing Power-On Self Test.	9-1
	Viewing Switch Status	9-3
	Viewing Port Information	9-4
	Viewing Equipment Status	9-8
	Viewing the System Message Log	9-10
	Viewing the Port Log	9-10
	Configuring for syslogd	9-12
	Configuring the Host	9-13
	Configuring the Switch.	9-13
	Viewing and Saving Diagnostic Information	9-14
	Setting Up Automatic Trace Dump Transfers	9-15
Chapter 10	Troubleshooting	
	Most Common Problem Areas	10-1
	Gathering Information for Technical Support.	10-2
	Troubleshooting Questions	10-2
	Analyzing Connection Problems	10-4
	Restoring a Segmented Fabric	10-8
	Correcting Zoning Setup Issues	10-10
	Recognizing MQ-WRITE Errors	10-12
	Correcting I2C Bus Errors	10-12
	Correcting Device Login Issues	10-14
	Identifying Media-Related Issues	10-16

Correcting Link Failures	10-19
Correcting Marginal Links	10-22
Inaccurate Information in the System Message Log	10-23
Recognizing the Port Initialization and FCP Auto Discovery Process.	10-24

Section II *Licensed Features*

Chapter 11 Administering Extended Fabrics

About Extended Link Buffer Allocation.	11-1
SilkWorm 3016, 3250, 3850, 3900, 12000, 24000, and 48000 (FC2-16 Port Blades)	11-1
SilkWorm 4012, 4100, and 48000 (FC4-16 and FC4-32 Port Blades). . .	11-1
Fabric Considerations	11-2
Choosing an Extended ISL Mode.	11-3
Configuring External Ports	11-4
Configuring an Extended ISL	11-7
Trunking Over Distance	11-8

Chapter 12 Administering ISL Trunking

Overview.	12-1
Standard Trunking Criteria	12-2
Fabric Considerations	12-3
Initializing Trunking on Ports.	12-4
Monitoring Traffic	12-4
Enabling and Disabling ISL Trunking	12-6
Setting Port Speeds	12-7
Displaying Trunking Information	12-8
Trunking Over Extended Fabrics	12-9
Troubleshooting Trunking Problems	12-9
Listing Link Characteristics	12-9
Recognizing Buffer Underallocation	12-10

Chapter 13 Administering Advanced Zoning

Zoning Terminology	13-1
Zoning Concepts	13-2
Zone Types	13-2
Zone Objects	13-5
Zone Aliases	13-5
Zone Configurations	13-5
Zoning Enforcement	13-6
Rules for Configuring Zones	13-10
Creating and Managing Zone Aliases	13-11
Creating and Maintaining Zones	13-14
.....	13-15
Merging Zones	13-16
Creating and Modifying Zoning Configurations	13-18
Maintaining Zone Objects	13-22
Managing Zoning Configurations in a Fabric	13-24
Adding a New Switch or Fabric	13-25
Splitting a Fabric	13-27
Using Zoning to Administer Security	13-27
Resolving Zone Conflicts	13-27

Chapter 14 Administering Advanced Performance Monitoring

Displaying and Clearing the CRC Error Count	14-3
Monitoring End-to-End Performance	14-3
Adding End-to-End Monitors	14-4
Setting a Mask for End-to-End Monitors	14-6
Deleting End-to-End Monitors	14-7
Monitoring Filter-Based Performance	14-7
Adding Standard Filter-based Monitors	14-8
Adding Custom Filter-Based Monitors	14-8
Deleting Filter-Based Monitors	14-10

Monitoring ISL Performance	14-10
Monitoring Trunks	14-10
Displaying Monitor Counters	14-11
Clearing Monitor Counters	14-14
Saving and Restoring Monitor Configurations	14-15
Collecting Performance Data	14-15

Appendix A Configuring the PID Format

About PIDs and PID Binding	A-1
Summary of PID Formats	A-1
Impact of Changing the Fabric PID Format	A-2
Host Reboots	A-2
Static PID Mapping Errors	A-3
Changes to Configuration Data	A-3
Selecting a PID format	A-4
Evaluating the Fabric	A-5
Planning the Update Procedure	A-7
Online Update	A-8
Offline Update	A-8
Hybrid Update	A-9
Changing to Core PID Format	A-9
Changing to Extended Edge PID Format	A-10
Converting Port Number to Area ID	A-12
Performing PID Format Changes	A-15
Basic Procedure	A-15
HP/UX Procedure	A-16
AIX Procedure	A-18
Swapping Port Area IDs	A-19

Appendix B Configuring Interoperability Mode

Vendor Switch Requirements	B-1
Brocade Switch Requirements	B-2

Supported Brocade Features	B-2
Unsupported Brocade Features	B-2
Configuration Recommendations	B-3
Configuration Restrictions	B-3
Zoning Restrictions	B-4
Zone Name Restrictions	B-4
Enabling and Disabling Interoperability Mode	B-5

Appendix C Using Remote Switch

Appendix D Understanding Legacy Password Behavior

Password Management Information	D-1
Password Prompting Behaviors	D-3
Password Migration During Firmware Changes	D-4
Password Recovery Options	D-5

Appendix E Zone Merging Scenarios

Appendix F Upgrading Firmware in Single CP Mode

Glossary

Index

About This Document

This document is a procedural guide to help SAN administrators configure and manage a storage area network (SAN) using the Brocade Fabric OS Command Line Interface (CLI).

How This Document Is Organized

The document is divided into two sections; the first, “Standard Features,” contains the following topics:

- [Chapter 1, “Introducing Fabric OS CLI Procedures”](#) gives a brief overview of Fabric OS, discusses the differences between SilkWorm switches and directors, and explains the Fabric OS CLI Help feature.
- [Chapter 2, “Performing Basic Configuration Tasks”](#) provides typical connection and configuration procedures.
- [Chapter 3, “Configuring Standard Security Features”](#) provides procedures for basic password and user account management.
- [Chapter 4, “Maintaining Configurations and Firmware”](#) provides configuration backup and firmware installation procedures.
- [Chapter 5, “Configuring SilkWorm 12000, 24000, and 48000 Directors”](#) provides information and procedures specific to SilkWorm 12000, 24000 and 48000 models. Because the SilkWorm 12000, 24000 and 48000 models have CP blades and port blades, they require procedures that are not relevant to the SilkWorm 3016, 3250, 3850, 3900, and 4100 fixed-port models.
- [Chapter 6, “Routing Traffic”](#) provides information and procedures for using switch routing features.
- [Chapter 7, “Administering FICON Fabrics”](#) provides procedures for use of the Brocade FICON Fabrics licensed feature.
- [Chapter 8, “Configuring the Distributed Management Server”](#) describes the use of the SAN management application.
- [Chapter 9, “Working With Diagnostic Features”](#) provides information about diagnostic and status-determining features, particularly system message logging.
- [Chapter 10, “Troubleshooting”](#) provides problem solving information and procedures.
- [Chapter 11, “Administering Extended Fabrics”](#) provides procedures for use of the Brocade Extended Fabrics licensed feature.

The second section, “Licensed Features,” contains the following topics:

- [Chapter 12, “Administering ISL Trunking”](#) provides procedures for use of the Brocade ISL Trunking licensed feature.

- [Chapter 13, “Administering Advanced Zoning”](#) provides procedures for use of the Brocade Advanced Zoning licensed feature.
- [Chapter 14, “Administering Advanced Performance Monitoring”](#) provides procedures for use of the Brocade Advanced Performance Monitoring licensed feature.
- [Appendix A, “Configuring the PID Format”](#) provides information about the various switch PID formats available and procedures for setting the PID format.
- [Appendix B, “Configuring Interoperability Mode”](#) provides information about using SilkWorm switches with other brands of switches.
- The other appendices provide special procedures or legacy information for previous Fabric OS versions.
- The glossary defines terms used in this document.
- The index provides links to the exact pages on which specific information is located.

Supported Hardware and Software

This document is specific to Fabric OS v5.0.1 running on the following Brocade SilkWorm product models:

- Brocade SilkWorm 200E switch
- Brocade SilkWorm 3014 switch
- Brocade SilkWorm 3016 switch
- Brocade SilkWorm 3250 switch
- Brocade SilkWorm 3850 switch
- Brocade SilkWorm 3900 switch
- Brocade SilkWorm 4012 switch
- Brocade SilkWorm 4100 switch
- Brocade SilkWorm 12000 director
- Brocade SilkWorm 24000 director
- Brocade SilkWorm 48000 director

In those instances in which procedures or parts of procedures documented here apply to some switches but not to others, this guide identifies which switches are supported and which are not.

This document sometimes mentions other Fabric OS releases to highlight the changes in the latest release or to point out interoperability issues with other SilkWorm models. It also specifies when procedures or steps of procedures apply only to specific SilkWorm models.

Although many different software and hardware configurations are tested and supported by Brocade Communications Systems, Inc. for this Brocade Fabric OS release, documenting all possible configurations and scenarios is beyond the scope of this document.

What's New in This Document

The following changes were made in the 53-0000518-07 edition of this document:

- The book has been reorganized; Section I contains standard Fabric OS features, and Section II contains licensed Fabric OS features.
- The RADIUS configuration procedures were revised.
- New features and procedures related to the SilkWorm 48000 switch have been added
- New features and procedures related to the SilkWorm 200E switch have been added.
- New features and procedures related to the SilkWorm 4012 switch have been added.

Document Conventions

This section describes text formatting conventions and important notice formats.

For readability, command names in the narrative portions of this guide are presented in mixed letter casing: for example, **switchShow**. In actual examples, command letter casing is often all lowercase. Otherwise, this manual specifically notes those cases in which a command is case sensitive.

Text Formatting

The narrative-text formatting conventions that are used in this document are as follows:

bold text	Identifies command names Identifies GUI elements Identifies keywords and operands Identifies text to enter at the GUI or CLI
<i>italic text</i>	Provides emphasis Identifies variables Identifies paths and Internet addresses Identifies document titles
code text	Identifies CLI output Identifies syntax examples

For readability, command names in the narrative portions of this guide are presented in mixed lettercase: for example, **switchShow**. In actual examples, command lettercase is often all lowercase. Otherwise, this manual specifically notes those cases in which a command is case sensitive. FICON commands are an exception to this convention.

Notes, Cautions, and Warnings

The following notices appear in this document.



Note

A note provides a tip, emphasizes important information, or provides a reference to related information.



Caution

A caution alerts you to potential damage to hardware, firmware, software, or data.



Warning

A warning alerts you to potential danger to personnel.

Additional Information

This section lists additional Brocade and industry-specific documentation that you might find helpful.

Brocade Resources

The following related documentation is provided on the Brocade Documentation CD-ROM and on the Brocade Web site, through Brocade Connect.



Note

Go to <http://www.brocade.com> and click **Brocade Connect** to register at no cost for a user ID and password.

Fabric OS

- *Fabric OS Command Reference Manual*
- *Fabric OS MIB Reference Manual*
- *Fabric OS System Error Message Reference Manual*

Fabric OS Optional Features

- *Web Tools Administrator's Guide*
- *Fabric Watch Administrator's Guide*
- *Secure Fabric OS Administrator's Guide*

SilkWorm 48000

- *SilkWorm 48000 Hardware Reference Manual*
- *SilkWorm 48000 QuickStart Guide*

SilkWorm 24000

- *SilkWorm 24000 Hardware Reference Manual*

- *SilkWorm 24000 QuickStart Guide*

SilkWorm 12000

- *SilkWorm 12000 Hardware Reference Manual*
- *SilkWorm 12000 QuickStart Guide*

SilkWorm 4100

- *SilkWorm 4100 Hardware Reference Manual (for v4.4.x and later software)*
- *SilkWorm 4100 QuickStart Guide (for v4.4.x and later software)*

SilkWorm 3900

- *SilkWorm 3900 Hardware Reference Manual (for v4.x software)*
- *SilkWorm 3900 QuickStart Guide (for v4.x software)*

SilkWorm 3250/3850

- *SilkWorm 3250/3850 Hardware Reference Manual (for v4.x software)*
- *SilkWorm 3250/3850 QuickStart Guide (for v4.x software)*

SilkWorm 3016

- *SilkWorm 3016 Hardware Reference Manual (for v4.2.x and later software)*
- *SilkWorm 3016 QuickStart Guide (for v4.2.x and later software)*
- *Brocade Enterprise and Entry SAN Switch Modules for IBM eServer BladeCenter Design, Deployment and Management Guide (DDM), which is available at this Web site:*
- <http://www-307.ibm.com/pc/support/site.wss/document.do?lnidocid=MIGR-55327>.

SilkWorm 3014

- *SilkWorm 3014 Hardware Reference Manual (for v5.x software)*
- *SilkWorm 3014 QuickStart Guide (for v5.x software)*

SilkWorm 200E

- *SilkWorm 200E Hardware Reference Manual*
- *SilkWorm 200E QuickStart Guide*

For practical discussions about SAN design, implementation, and maintenance, you can obtain *Building SANs with Brocade Fabric Switches* through:

<http://www.amazon.com>

For information about how to use Fabric OS features in a SAN solution, refer to the *Brocade SilkWorm Design, Deployment, and Management Guide*. You can obtain this guide through the Brocade Connect Web site:

<http://www.brocadeconnect.com>

For additional Brocade documentation, visit the Brocade SAN Info Center and click the Resource Library location:

<http://www.brocade.com>

Release notes are available on the Brocade Connect Web site and are also bundled with the Fabric OS firmware.

Other Industry Resources

For additional resource information, visit the Technical Committee T11 Web site. This Web site provides interface standards for high-performance and mass storage applications for Fibre Channel, storage management, as well as other applications:

<http://www.t11.org>

For information about the Fibre Channel industry, visit the Fibre Channel Industry Association Web site:

<http://www.fibrechannel.org>

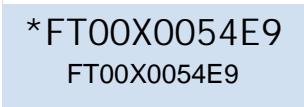
Getting Technical Help

Contact your switch support supplier for hardware, firmware, and software support, including product repairs and part ordering. To expedite your call, have the following information available:

1. General Information
 - Technical Support contract number, if applicable
 - Switch model
 - Switch operating system version
 - Error messages received
 - **supportSave** command output
 - Serial console and/or telnet session logs
 - syslog message logs
 - Detailed description of the problem and specific questions
 - Description of any troubleshooting steps already performed and results

2. Switch Serial Number

The switch serial number and corresponding bar code are provided on the serial number label, as illustrated below.



*FT00X0054E9
FT00X0054E9

The serial number label is located as follows:

- SilkWorm 3014 switches: Top of chassis, under the insertion arm
- SilkWorm 3016 and 4012 switches: Side of switch module
- SilkWorm 200E, 3200, and 3800 switches: Nonport side of chassis
- SilkWorm 3250, 3850, and 3900 switches: Bottom of chassis
- SilkWorm 4100 switches: On the switch ID pull-out tab located on the port side and on the inside of the chassis, near power supply # 1 (the power supply on the right when looking at the nonport side of the switch).
- SilkWorm 12000, 24000, and 48000 directors: Inside front of chassis, on wall to left of ports

3. World Wide Name (WWN)

- *SilkWorm 200E, 3014, 3016, 3250, 3850, 3900, 4012, and 4100 switches and SilkWorm 12000, 24000 and 48000 directors:* Provide the license ID. Use the **licenseIdShow** command to display the license ID.
- *All other SilkWorm switches:* Provide the switch WWN. Use the **wwn** command to display the switch WWN.

Document Feedback

Brocade Communications Systems, Inc. makes every effort to ensure the accuracy and completeness of this document. However, if you find an error or an omission, or you think that a topic needs further development, we want to hear from you. Forward your feedback to documentation@brocade.com. Provide the title and version number and as much detail as possible about your issue, including the topic heading and page number and your suggestions for improvement.

Standard Features

This section describes standard Brocade Fabric OS features, and includes the following chapters:

- [Chapter 1, *Introducing Fabric OS CLI Procedures*](#)
- [Chapter 2, *Performing Basic Configuration Tasks*](#)
- [Chapter 3, *Configuring Standard Security Features*](#)
- [Chapter 4, *Maintaining Configurations and Firmware*](#)
- [Chapter 5, *Configuring SilkWorm 12000, 24000, and 48000 Directors*](#)
- [Chapter 6, *Routing Traffic*](#)
- [Chapter 7, *Administering FICON Fabrics*](#)
- [Chapter 8, *Configuring the Distributed Management Server*](#)
- [Chapter 9, *Working With Diagnostic Features*](#)
- [Chapter 10, *Troubleshooting*](#)

Introducing Fabric OS CLI Procedures

This guide contains procedures for configuring and managing a Brocade® SilkWorm® Storage Area Network (SAN) using the Fabric OS Command Line Interface (CLI).

The guide applies to the following Brocade product models:

- SilkWorm switches: 200E, 3016, 3250, 3850, 3900, 4012, and 4100

These SilkWorm models contain a fixed number of ports (they are *fixed-port switches*). The SilkWorm 4100 and 200E models allow you to license and activate extra fixed ports with the Ports on Demand feature.

- SilkWorm directors: 12000, 24000, and 48000

These SilkWorm models can contain a variable number of ports, which you install by plugging port blades into the director chassis. The SilkWorm 48000 can have up to 256 ports, and the SilkWorm 12000 and 24000 can have up to 128 ports.

About Procedural Differences

As a result of the differences between fixed-port and variable-port devices, procedures sometimes differ among SilkWorm models. Also, because the domain architecture of the SilkWorm 12000 differs from that of the SilkWorm 24000 and SilkWorm 48000, there are sometimes procedural differences among these models. As new SilkWorm models are introduced, new features sometimes apply only to those models.

When procedures or parts of procedures apply to some models but not others, this guide identifies the specifics for each model. For example, a number of procedures that apply only to variable-port devices are found in [“Configuring SilkWorm 12000, 24000, and 48000 Directors” on page 5-1](#). Procedures that apply only to the SilkWorm 3016 or 4100 model are labeled as such.



Note

When command examples in this guide show user input enclosed in quotation marks, the quotation marks are required for versions earlier than v4.0.0. They are optional in later versions, unless specifically called for in the procedures.

Scope and References

Although many different software and hardware configurations are tested and supported by Brocade Communications Systems, Inc., documenting all possible configurations and scenarios is beyond the scope of this document. In some cases, earlier releases are highlighted to present considerations for interoperating with them.

The hardware reference manuals for SilkWorm products describe how to power up devices and set their IP addresses. After the IP address is set, you can use the CLI procedures contained in this guide.

This guide provides only the level of detail required to perform the procedures. If you need more information about the commands used in the procedures, refer to online help or to the *Fabric OS Command Reference Manual*.

There are several access methods that you can use to configure a switch. These are listed with their respective documents:

- Command line interface (CLI)
 - A telnet session into logical switches
 - A telnet session into active and standby CPs for director class switches
 - A serial console, including active and standby CPs for director class switches
 - An optional modem, which behaves like a serial console port

For CLI details, refer to the *Fabric OS Command Reference Manual*.

- Web Tools

For Web Tools procedures, refer to the *Web Tools Administrator's Guide*.

- Fabric Manager

For Fabric Manager procedures, refer to the *Fabric Manager Administrator's Guide*.

- A third-party application using the API

For third-party application procedures, refer to the third-party API documentation.

About the CLI

Fabric OS CLI is the complete fabric management tool for Brocade SANs that enables you to:

- Access the full range of Fabric OS features, based on license keys.
- Configure, monitor, dynamically provision, and manage every aspect of the SAN.
- Configure and manage the Brocade fabric on multiple efficient levels.
- Identify, isolate, and manage SAN events across every switch in the fabric.
- Manage switch licenses.
- Perform fabric stamping.

To manage a switch using telnet, SNMP, and Web Tools, the switch must be connected to a network through the switch Ethernet port (out of band) or from the Fibre Channel (in band). The switch must be configured with an IP address to allow for the network connection. Refer to the hardware manual for your specific switch for information on physically connecting to the switch.

You can access switches from different connections, such as Web Tools, CLI, and API. When these connections are simultaneous, changes from one connection might not be updated to the other, and some modifications might be lost. When simultaneous connections are used, make sure that you do not overwrite the work of another connection.

In a mixed fabric containing switches running various Fabric OS versions, you should use the latest-model switches running the most recent release for the primary management tasks. The principal management access should be set to the core switches in the fabric. For example, to run Secure Fabric OS®, use the latest-model switch as the primary FCS, the location to perform zoning tasks, and the time server.

A number of management tasks are designed to make fabric-level changes; for example, zoning commands make changes that affect the entire fabric. When executing fabric-level configuration tasks, allow time for the changes to propagate across the fabric before executing any subsequent tasks. For a large fabric, it might take a few minutes.

Help Information

Each Fabric OS command provides Help information that explains the command function, its possible operands, its level in the command hierarchy, and additional pertinent information.

Displaying Command Help

To display help information:

1. Connect to the switch and log in as admin.
2. To display a list of all command help topics for a given login level, enter the **help** command with no arguments. For example, if you are logged in as *user* and enter the **help** command, a list of all *user*-level commands that can be executed will be listed. The same rule applies to the *admin* role. In addition, any user-configured command that uses a *switchAdmin* role will also see commands available to users with the *switchAdmin* role.
3. To display help for a specific command for a given login level, enter **help command**, where *command* is the name of the command for which you need information.

Example

```
switch:admin> help configure
Administrative Commands          configure(lm)
NAME
    configure - change system configuration settings
SYNOPSIS
    configure
AVAILABILITY
    admin
DESCRIPTION
    This command changes some system configuration settings,
    including:
    o Arbitrated loop settings
    o Switch fabric settings
    o System services settings
    o Virtual channel settings
    (output truncated)
```

Displaying Additional Help Topics

The following commands provide help files for specific topics:

diagHelp	Diagnostic help information
ficonHelp	FICON help information
fwHelp	Fabric Watch help information
licenseHelp	License help information
perfHelp	Performance Monitoring help information
routeHelp	Routing help information
trackChangesHelp	Track Changes help information
zoneHelp	Zoning help information

Performing Basic Configuration Tasks

This chapter contains procedures for performing basic switch configuration tasks using the Fabric OS command line interface (CLI).

Connecting to the Command Line Interface

You can connect to the command line interface (CLI) either through a telnet connection or through the serial port.

To connect with telnet

1. Verify that the switch is connected to the IP network through the RJ-45 Ethernet port.

Switches in the fabric that are not connected via Ethernet can be managed through switches that are using IP over Fibre Channel. The embedded port must have an assigned IP address.

SilkWorm 3016: Ethernet connectivity is facilitated through the IBM eServer BladeCenter Management Module.

2. Open a telnet connection using the IP address of the logical switch to which you want to connect. If you telnet to the active CP or log in to the active CP console, you are prompted for the switch number when the platform is set up in dual (or multiple) switch mode. For example, SilkWorm 24000 will not prompt you if you are using configuration option 1, but will prompt you if you have used configuration options 2-4. Refer to [Chapter 5, “Configuring SilkWorm 12000, 24000, and 48000 Directors”](#) for details about director configuration options.

SilkWorm 12000 and SilkWorm 24000 (configured with two domains) have two logical switches (**sw0** and **sw1**).

The login prompt is displayed when the telnet connection finds the switch in the network.

3. Enter the account ID (defaults are user or admin) at the login prompt.
4. Enter the password.

The default password is: password

If you have not changed the system passwords from the default, you are prompted to change them. Enter the new system passwords, or press **Ctrl-c** to skip the password prompts.

5. Verify that the login was successful. The prompt displays the switch name and user ID to which you are connected.

```
login: admin
password: xxxxxxxx
switch:admin>
```

Observe these considerations for telnet connections:

- Never change the IP address of the switch while two telnet sessions are active; if you do, your next attempt to log in fails. To recover, gain access to the switch by one of these methods:
 - You can use Web Tools and perform a fast boot. When the switch comes up, the telnet quota is cleared. (For instructions on performing a fast boot with Web Tools, refer to the *Web Tools Administrator's Guide*.)
 - If you have the required privileges, you can connect through the serial port, log in as root, and use operating system commands to identify and kill the telnet processes without disrupting the fabric.
- For admin level accounts, Fabric OS limits the number of simultaneous telnet sessions per switch to two. For more details on session limits, refer to [“Configuring the Telnet Interface” on page 3-3](#) and [“Creating and Maintaining User-Defined Accounts” on page 3-6](#).

To connect through the serial port

1. Connect the serial cable to the serial port on the switch and to an RS-232 serial port on the workstation.

If the serial port on the workstation is RJ-45 instead of RS-232, remove the adapter on the end of the serial cable and insert the exposed RJ-45 connector into the RJ-45 serial port on the workstation.

2. Open a terminal emulator application (such as HyperTerminal on a PC, or TERM, TIP, or Kermit in a UNIX environment), and configure the application as follows:

- In a Windows environment:

Parameter	Value
Bits per second	9600
Databits	8
Parity	None
Stop bits	1
Flow control	None

- In a UNIX environment, enter the following string at the prompt: **`tip /dev/ttyb -9600`**.

If ttyb is already in use, you can use ttya (enter **`tip /dev/ttya -9600`**).

Observe these considerations for serial connections:

- Some procedures require that you connect through the serial port; for example, setting the IP address or setting the boot PROM password.
- If secure mode is enabled, connect through the serial port of the primary FCS switch.
- **SilkWorm 3016:** This model does not have an external serial port.
- **SilkWorm 12000, 24000, and 48000:** You can connect to CP0 or CP1 using either of the two serial ports.

Setting the IP Address

You must connect through the serial port to set the IP address (refer to [“To connect through the serial port” on page 2-2](#)). After connecting, use the **ipAddrSet** command to set the IP address.

To set the IP address on the SilkWorm 3016

The SilkWorm 3016 does not have an Ethernet port or an external serial port. You must configure the IP address for this model through the IBM eServer BladeCenter management module. Refer to the *Brocade Enterprise and Entry SAN Switch Modules for IBM eServer BladeCenter Design, Deployment and Management Guide DDM* for the procedure.



Caution

The use of IP address 0.0.0.0 is not supported. Do not use this address.

Fabric OS beginning with v2.6.0, v3.1.0, and v4.0.0 supports Classless Inter-Domain Routing (CIDR).

Setting the Default Account Passwords

For each logical switch (domain), there are admin and user default access accounts. These accounts designate different levels of authorization—called *roles*—for using the system:

- *admin* level for administrative use
- *user* level for nonadministrative use, such as monitoring system activity
- *switchAdmin* level for administrative use, except for security, user management, and zoning

Two accounts—factory and root—are reserved for development and manufacturing purposes. You can change their passwords, which is optional, but you should not use these accounts under normal circumstances.

[Table 2-1](#) shows the number of simultaneous login sessions allowed for each role.

Table 2-1 Maximum Number of Simultaneous Sessions

User Name	Maximum Sessions
admin	2
user	4

SilkWorm 3016: The default administrative account is called USERID (uppercase characters). For instructions on changing this name, refer to [“To change the default administrative account name on the SilkWorm 3016” on page 2-5](#).

SilkWorm 200E, 3016, 3250, 3850, 3900, 4012, and 4100 switches and SilkWorm 24000 and 48000 (default configuration with one domain): There is one set of default access accounts.

SilkWorm 12000 and 24000 (configured with two domains): Each logical switch has its own set of default access accounts. The default account names and passwords are the same for both of the logical switches.

You can also create up to 15 additional accounts per logical switch and designate their roles as either admin, switchAdmin, or user. Refer to the procedures for doing so in [“Creating and Maintaining User-Defined Accounts” on page 3-6](#).

For large enterprises, Fabric OS supports RADIUS services, as described in [“Setting Up RADIUS AAA Service” on page 3-9](#).

In addition to the account access passwords, each switch can set a boot PROM password. For greater security, it is recommended that you set this password to protect system boot parameters from unauthorized access. Refer to [“Setting the Boot PROM Password” on page 3-34](#).)

Each of the default access accounts has an associated password. The first time you connect to a Fabric OS switch you are prompted to change these default account passwords.

If you do not change the default passwords, you are prompted to do so at each subsequent login until all system passwords have been changed from the default values. Thereafter, use the **passwd** command to change passwords.

For more background information on passwords, refer to [“Changing an Account Password” on page 3-8](#).

To change the default passwords at login

1. Connect to the switch and log in as admin.

SilkWorm 3016: The default password for all default accounts is: PASSWORD (uppercase characters, and the 0 is the number 0, not the alphabetic character “O.”)

Other models: The default password for all default accounts is: password

2. At each of the “Enter new password” prompts, either enter a new password or skip the prompt.

You can skip a prompt by pressing **Enter**. You can bypass all further prompts by using **Ctrl-c**.

Although the root and factory accounts are not meant for general use, you should change their passwords if prompted to do so and save the passwords in case they are needed for recovery purposes.

Passwords can be from 8 to 40 characters long. They must begin with an alphabetic character. They can include numeric characters, the dot (.), and the underscore (_). They are case-sensitive, and they are not displayed when you enter them on the command line.

You cannot reuse the default passwords.



Note

Record the passwords exactly as entered and store them in a secure place, because recovering passwords requires significant effort and fabric downtime.

The initial login prompt accepts a maximum password length of eight characters. Any characters beyond the eighth character are ignored. Only the default password is subject to the eight character limit. Any password set by the user can have a length from 8 to 40 characters.

```

login: admin
Password:
Please change your passwords now.
Use Control-C to exit or press 'Enter' key to proceed.
for user - root
Changing password for root
Enter new password: *****
Password changed.
Saving password to stable storage.
Password saved to stable storage successfully.
Please change your passwords now.
for user - factory
Changing password for factory
Enter new password: *****
Password changed.
Saving password to stable storage.
Password saved to stable storage successfully.
Please change your passwords now.
for user - admin
Changing password for admin
Enter new password: *****
Password changed.
Saving password to stable storage.
Password saved to stable storage successfully.
Please change your passwords now.
for user - user
Changing password for user
Enter new password: *****
Password changed.
Saving password to stable storage.
Password saved to stable storage successfully.
switch:admin>

```

To change the default administrative account name on the SilkWorm 3016

Use the **userRename** command to rename the administrative account from USERID to the Brocade default, admin, before enabling security; otherwise, the SilkWorm 3016 switch will not be allowed in the secure fabric. If you previously renamed the default user-level account from user to some other name, rename it to the default, user, before enabling security.

```
brocadessm:USERID> userrename "USERID" "admin"
```

After you rename the account, at next login the system prompt changes to:

```
brocadessm:admin>
```

Setting the Date and Time

Switches maintain the current date and time in flash memory. Date and time are used for logging events. Switch operation does not depend on the date and time; a switch with an incorrect date and time value still functions properly. However, because the date and time are used for logging, you should set them correctly.

**Note**

If secure mode is not enabled, a change in date or time to one switch is forwarded to the principal switch and distributed to the fabric. If secure mode is enabled, date or time changes can be made only on the primary FCS switch and distributed to the fabric.

To set the date and time

1. Connect to the switch and log in as admin.
2. Enter the **date** command at the command line, using the following syntax:

```
date "mmddHHMMyy"
```

The values represent the following:

- mm is the month; valid values are 01 through 12.
- dd is the date; valid values are 01 through 31.
- HH is the hour; valid values are 00 through 23.
- MM is minutes; valid values are 00 through 59.
- yy is the year; valid values are 00 through 99 (values greater than 69 are interpreted as 1970 through 1999, and values less than 70 are interpreted as 2000-2069).

Example

```
switch:admin> date
Fri Jan 29 17:01:48 UTC 2000
switch:admin> date "0227123003"
Thu Feb 27 12:30:00 UTC 2003
switch:admin>
```

For details about how to change time zones, refer to **tsTimeZone** command in the *Fabric OS Command Reference Manual*.

You can synchronize the local time of the Principal or Primary Fabric Configuration Server (FCS) switch to an external NTP server.

To synchronize local time with an external source

1. Connect to the switch and log in as admin.
2. Enter the following command:

```
tsclockserver ipaddr
```

where *ipaddr* is the IP address of the NTP server, which the switch must be able to access. This operand is optional; by default this value is LOCL, which uses the local clock of the principal or primary switch as the clock server.

Example

```
switch:admin> tsclockserver
LOCL
switch:admin> tsclockserver "132.163.135.131"
switch:admin> tsclockserver
132.163.135.131
switch:admin>
```

It is suggested that you synchronize time with an external NTP server, as described on [page 2-6](#). If you cannot do so, use the next procedure.

To set the time zone

1. Connect to the switch and log in as admin.
2. Enter the **tsTimeZone** command as follows:

tsTimeZone [houroffset [, minuteoffset]]

- For Pacific Standard Time enter **tsTimeZone -8,0**
- For Central Standard Time enter **tsTimeZone -6,0**
- For Eastern Standard Time enter **tsTimeZone -5,0**

The default time zone for switches is Universal Time Conversion (UTC), which is 8 hours ahead of Pacific Standard Time (PST). For additional time zone conversions, refer to [Table 2-2 on page 2-7](#).

The parameters do not apply if the time zone of the switch has already been changed from the default (8 hours ahead of PST).

Refer to the **tsTimeZone** command in the *Fabric OS Command Reference Manual* for more detailed information about the command parameters.

Repeat the procedure on all switches for which the Time Zone needs to be set.

This only needs to be done one time, because the value is written to nonvolatile memory.

For U.S. time zones, use [Table 2-2](#) to determine the correct parameter for the **tsTimeZone** command.

Table 2-2 Conversion from UTC to Local Time

Local Time	ts timezone Parameter (Difference From Universal Time Conversion)
Atlantic Standard	-4,0
Atlantic Daylight	-3,0
Eastern Standard	-5,0
Eastern Daylight	-4,0
Central Standard	-6,0
Central Daylight	-5,0
Mountain Standard	-7,0
Mountain Daylight	-6,0
Pacific Standard	-8,0

Table 2-2 Conversion from UTC to Local Time (Continued)

Local Time	tstimezone Parameter (Difference From Universal Time Conversion)
Pacific Daylight	-7,0
Alaskan Standard	-9,0
Alaskan Daylight	-8,0
Hawaiian Standard	-10,0

Maintaining Licensed Features

Feature licenses might be part the licensed Paper Pack supplied with switch software, or you can purchase licenses separately from your switch vendor, who will provide you with keys to unlock the features. License keys are provided on a per-chassis basis, so for products that support multiple logical switches (domains), a license key applies to all domains within the chassis.

To unlock a licensed feature, you can either use the license key provided in the Paper Pack supplied with switch software, or follow this procedure to generate a license key at the Brocade Web site (<http://www.brocade.com>).



Note

You need the following items for each chassis to be licensed:

- transaction key

This key is in the Paper Pack supplied with switch software. Or, when you purchase a license your switch vendor gives you a transaction key to be used for obtaining a license key.

- License ID

To see a switch license ID, use the **licenseIdShow** command.

To unlock a licensed feature

1. If you already have a license key, go to [step 10](#).

If you do not have a key, launch an Internet browser and go to the Brocade Web site at <http://www.brocade.com>.

2. Click **products**.

3. Click **Software Products**.

4. In the **Related Links** panel on the right side of the page, select **Software License Keys**. The Software License Keys instruction page appears.

5. If you want to generate a single license key, select **Generate 1 license key**.

If you want to generate multiple license keys, select **Batch Generation of Licenses**.

The Software License Key instruction page appears.

6. Enter the requested information in the required fields.

When generating multiple license keys, enter the worldwide names and transaction keys in the table at the bottom of the screen. If you need additional rows in the table, select **Add More Rows**.

7. Click **Next**.

A verification screen appears.

8. Verify that the information appears correctly.

Click **Submit** if the information displayed is correct. If the information is incorrect, Click **Previous** and change the information.

9. After the information is corrected, click **Submit**.

An information screen displays the license keys.

You also receive an e-mail from Brocade with the keys and installation instructions.

10. Activate and verify the license as follows:

- a. Connect to the switch and log in as admin.
- b. Activate the license using the **licenseAdd** command.

```
switch:admin> licenseadd "key"
```

The license key is case sensitive and must be entered exactly as given. The quotation marks are optional.

For SilkWorm 12000, 24000, and 48000 models, licenses are effective on both CP blades and on all logical switches, but are valid only when the CP blade is inserted into a chassis that has an appropriate license ID stored in the WWN card. If a CP is moved from one chassis to another, the license works in the new chassis only if the WWN card is the same in the new chassis. Otherwise, a new license key is generated.

For example, if you swap one CP blade at a time, or replace a single CP blade, then the existing CP blade (the active CP blade) propagates the licenses to the new CP blade.

If you move a standby CP from one chassis to another, then the active CP will propagate its configuration (including license keys).

- c. Verify that the license was added by entering the **licenseShow** command. The licensed features currently installed on the switch are listed. If the feature is not listed, enter the **licenseAdd** command again.
- d. Some features may require additional configuration, or you might need to disable and reenable the switch to make them operational; refer to the feature documentation for details.

Example

```
switch:admin> licenseshow
SbeSdQdQySyrTeJ:
Web license
Zoning license
Fabric license
Remote Switch license
Extended Fabric license
Fabric Watch license
Performance Monitor license
Trunking license
Security license
SbbebdQS9QTscfcB:
Ports on Demand license - additional 8 port upgrade
SbbebdQS9QTcgfcz:
Ports on Demand license - additional 8 port upgrade
```

To remove a licensed feature

1. Connect to the switch and log in as admin.
2. Enter the **licenseShow** command to display the active licenses.
3. Remove the license key using the **licenseRemove** command.

```
switch:admin> licenseremove "key"
```

The license key is case sensitive and must be entered exactly as given. The quotation marks are optional. After removing a license key, the optionally licensed feature is disabled when the switch is rebooted or when a switch disable or enable is performed.

4. Enter the **licenseShow** command to verify that the license is disabled.

```
switch:admin> licenseshow
bQebzbRdScRfc0iK:
  Web license
  Zoning license
SybbzQQ9edTzcc0X:
  Fabric license
switch:admin> licenseremove "bQebzbRdScRfc0iK"
removing license key "bQebzbRdScRfc0iK"
switch:admin>
```

After a reboot (or **switchDisable** and **switchEnable**):

```
switch:admin> licenseshow
SybbzQQ9edTzcc0X:
  Fabric license
switch:admin>
```

If there are no license keys, **licenseShow** displays “No licenses.”

Customizing the Switch Name

Switches can be identified by IP address, Domain ID, World Wide Name (WWN), or by customized switch names that are unique and meaningful.

Version 4.0.0 (and later) Switch names can be from 1 to 15 characters long, must begin with a letter, and can contain letters, numbers, or the underscore character. It is not necessary to use quotation marks.

The default names are:

- **SilkWorm 3016:** brocadessm
- **SilkWorm 200E, 3250, 3850, 3900, 4012, and 4100 switches:** swd77
- **SilkWorm 12000:** varies depending on number of logical switches
The two logical switches have different default names. The name “swd77” is used for the logical switch containing the port blades in slots 1 through 4, and “swd76” is used for the logical switch containing the port blades in slots 7 through 10.

- **SilkWorm 24000 and 48000:** swd77

**Note**

Changing the switch name causes a domain address format RSCN to be issued.

To customize the switch name

1. **SilkWorm 200E, 3016, 3250, 3850, 3900, 4012, and 4100 switches:** Proceed to the next step.
SilkWorm 12000, 24000, and 48000 directors: Open a telnet window for each logical switch and enter the **switchName** command.
2. Connect to the switch and log in as admin.
3. **SilkWorm 200E, 3016, 3250, 3850, 3900, 3012, and 4100 switches:** Proceed to the next step.
SilkWorm 24000 and 48000: If configured for one domain (the default) proceed to the next step. If configured with two domains, proceed as for the SilkWorm 12000.
SilkWorm 12000: Choose the logical switch that you want to change. Enter the value that corresponds to that logical region:
 - Enter 0 to configure logical switch 0 (slot 1 through 4)
 - Enter 1 to configure logical switch 1 (slot 7 through 10)
4. Enter the **switchName** command at the command line, using the following syntax:

```
switchname "newname"
```

where *newname* is the new name for the switch.
5. Record the new switch name for future reference.
6. **SilkWorm 12000, 24000, and 48000 configured with two domains:** Disconnect from the session and repeat the procedure for the second logical switch.

Example

```
switch:admin> switchname "switch62"  
Committing configuration...  
Done.  
switch62:admin>
```

Customizing the Chassis Name

Beginning with Fabric OS v4.4.0, it is recommended that you customize the chassis name for each switch. Some system logs identify switches by chassis names, so if you assign meaningful chassis names in addition to meaningful switch names, logs will be more useful.

To change the chassis name

1. Connect to the switch and log in as admin.
2. Enter the **chassisName** command at the command line, using the following syntax:

```
chassisname "newname"
```

where *newname* is the new name for the chassis.

Chassis names can be from 1 to 15 characters long, must begin with a letter, and can contain letters, numbers, or the underscore character. It is not necessary to use the quotation marks.

3. Record the new chassis name for future reference.

Disabling and Enabling a Switch

By default, the switch is enabled after power is applied and diagnostics and switch initialization routines have finished. You can disable and reenable it as necessary.

To disable a switch

1. Connect to the switch and log in as admin.
2. Enter the **switchDisable** command at the command line.

All Fibre Channel ports on the switch are taken offline. If the switch was part of a fabric, the fabric reconfigures.

To enable a switch

1. Connect to the switch and log in as admin.
2. Enter the **switchEnable** command at the command line.

All Fibre Channel ports that passed the POST test are enabled. If the switch has interswitch links to a fabric, it joins the fabric.

Disabling and Enabling a Port

All licensed ports are enabled by default. You can disable and reenable them as necessary. Ports that you activate with Ports on Demand must be enabled explicitly, as described in [“Activating Ports on Demand” on page 2-13](#).

To disable a port

1. Connect to the switch and log in as admin.
2. **SilkWorm 200E, 3016, 3250, 3850, 3900, 4012, and 4100 switches:** Enter the following command:

```
portdisable portnumber
```

where *portnumber* is the port number of the port you want to disable.

SilkWorm 12000, 24000, and 48000 directors: Enter the following command:

```
portdisable slotnumber/portnumber
```

where *slotnumber* and *portnumber* are the slot and port numbers of the port you want to disable.

If the port is connected to another switch, the fabric might reconfigure.

To enable a port

1. Connect to the switch and log in as admin.
2. **SilkWorm 200E, 3016, 3250, 3850, 3900, 4012, and 4100 switches:** Enter the following command:

```
portenable portnumber
```

where *portnumber* is the port number of the port you want to enable.

SilkWorm 12000, 24000, and 48000 directors: Enter the following command:

```
portenable slotnumber/portnumber
```

where *slotnumber* and *portnumber* are the slot and port numbers of the port you want to enable. (Slots are numbered 1 through 4 and 7 through 10, counting from left to right.)

If the port is connected to another switch, the fabric might reconfigure. If the port is connected to one or more devices, these devices become available to the fabric.

If you change port configurations during a switch failover, the ports might become disabled. To bring the ports online, reissue the **portEnable** command after the failover is complete.

Activating Ports on Demand

The SilkWorm 4100 model can be purchased with 16, 24, or 32 licensed ports. As your needs increase, you can activate unlicensed ports (up to the maximum of 32 ports) by purchasing and installing the Brocade Ports on Demand optional licensed product.

The SilkWorm 200E model can be purchased with 8 ports and no E_Port, 8 ports with full fabric access (SilkWorm 210E), and 16 ports with full fabric access (SilkWorm 240E). If you purchase the SilkWorm 200E with 8 ports enabled, you can activate unlicensed ports in 4-port increments up to 16 ports by purchasing and installing the Brocade Ports on Demand optional licensed product. You can also purchase an E_Port license upgrade if your switch does not yet support full fabric access.

Ports on Demand is ready to be unlocked in the switch firmware. Its license key might be part of the licensed Paper Pack supplied with switch software, or you can purchase the license key separately from your switch vendor. You might need to generate a license key from a transaction key supplied with your purchase. If so, launch an Internet browser and go to the Brocade Web site at <http://www.brocade.com>. Click **products**> **Software Products**> **Software License Keys** and follow the instructions to generate the key.

By default, ports 0 through 15 are activated on the SilkWorm 4100. Each Ports on Demand license activates the next group of eight ports, in numerical order. Before installing a license key, you must insert transceivers in the ports to be activated. Remember to insert the transceivers in the lowest group of inactive port numbers first. For example, if only 16 ports are currently active and you are installing one Ports on Demand license key, make sure to insert the transceivers in ports 16 through 23. If you later install a second license key, insert the transceivers in ports 24 through 31. (For details on inserting transceivers, refer to the *SilkWorm 4100 Hardware Reference Manual*).

After you install a license key, you must enable the ports to complete their activation. You can do so without disrupting switch operation by using the **portEnable** command on each port. Alternatively, you can disable and reenable the switch to activate ports.

To activate Ports on Demand

1. Connect to the switch and log in as admin.
2. Optionally, to verify the current states of the ports, use the **portShow** command.
In the **portShow** output, the Licensed field indicates whether the port is licensed or not.
3. Install the Brocade Ports on Demand license.
For instructions, refer to “[Maintaining Licensed Features](#)” on page 2-8.
4. Use the **portEnable** command to enable the ports.
5. Optionally, use the **portShow** command to check the newly activated ports.

If you remove a Ports on Demand license, the licensed ports will become disabled after the next platform reboot or the next port deactivation.

Making Basic Connections

You can make basic connections to devices and to other switches.

Before connecting a v4.0.0 (or later) switch to a fabric that contains switches running earlier firmware versions, you must first set the same PID format on all the switches. The presence of different PID formats in a fabric causes fabric segmentation.

For information on PID formats and related procedures, refer to “[Selecting a PID format](#)” on page A-4.

For information on configuring the routing of connections, refer to “[Routing Traffic](#)” on page 6-1.

For information on configuring extended interswitch connections, refer to “[Administering Extended Fabrics](#)” on page 11-1.

Connecting to Devices

To minimize port logins, power off all devices before connecting them to the switch. For devices that cannot be powered off, first use the **portDisable** command to disable the port on the switch, and then connect the device. When powering the devices back on, wait for each device to complete the fabric login before powering on the next one.

Connecting to Other Switches

Refer to the hardware user's guide of your specific switch for interswitch link (ISL) connection and cable management information. The standard (default) ISL mode is L0, which you can configure with the **portCfgLongDistance** command. ISL Mode L0 is a static mode, with the following maximum ISL distances:

- 10 km at 1 Gbit/second
- 5 km at 2 Gbit/second
- 2.5 km at 4 Gbit/second

ISL mode L0 is available on all Fabric OS releases. When you upgrade from Fabric OS v4.0.0 to Fabric 4.1.0 or later, all extended ISL ports are set automatically to L0 mode.

For information on extended ISL modes, which enable longer distance interswitch links, refer to [Chapter 11, “Administering Extended Fabrics”](#).

Working With Domain IDs

Although domain IDs are assigned dynamically when a switch is enabled, you can reset them manually so that you can control the ID number or to resolve a domain ID conflict when you merge fabrics.

If a switch already has a domain ID when it is enabled, and that domain ID conflicts with a switch already in the fabric, the conflict is automatically resolved. The process can take several seconds, during which time traffic is delayed.

The default domain ID for SilkWorm switches is 1.

The default domain ID applies to the logical switches in SilkWorm 12000, 24000, and 48000 directors that are configured for two domains. To prevent domain conflict, you can either disable one of the switches until the other is connected to the fabric, then reenabling the switches so that unique domain IDs are automatically assigned, or you can use the procedure [“To set the domain ID” on page 2-16](#) to make the domain IDs unique before connecting the logical switches to the fabric.



Caution

On switches running Fabric OS v4.0.0 and later, do not use domain ID 0, which is reserved for another purpose. The use of this domain ID can cause the switch to reboot continuously.

Avoid changing the domain ID on the FCS in secure mode. To minimize down time, change the domain IDs on the other switches in the secure fabric.

To display domain IDs

1. Connect to a switch and log in as admin.
2. Enter the **fabricShow** command.

Fabric information is displayed, including the domain ID (D_ID):

```

ras045:admin> fabricshow
Switch ID      Worldwide Name      Enet IP Addr      FC IP Addr      Name
-----
 1: fffc01 10:00:00:60:69:e4:00:3c 10.32.220.80      0.0.0.0          "ras080"
 2: fffc02 10:00:00:60:69:e0:01:46 10.32.220.1       0.0.0.0          "ras001"
 3: fffc03 10:00:00:60:69:e0:01:47 10.32.220.2       0.0.0.0          "ras002"
 5: fffc05 10:00:00:05:1e:34:01:bd 10.32.220.5       0.0.0.0          "ras005"
 6: fffc06 10:00:00:05:1e:34:02:3e 10.32.220.6       0.0.0.0          "ras006"
 7: fffc07 10:00:00:60:69:34:02:0c 10.32.220.7       0.0.0.0          "ras007"
10: fffc0a 10:00:00:60:69:80:04:46 10.32.220.10      10.32.219.0      "ras010"
11: fffc0b 10:00:00:60:69:80:04:47 10.32.220.11      10.32.219.1      "ras011"
15: fffc0f 10:00:00:60:69:80:47:74 10.32.220.15      0.0.0.0          "ras015"
16: fffc10 10:00:00:60:69:80:47:75 10.32.220.16      0.0.0.0          "ras016"
19: fffc13 10:00:00:05:1e:34:00:ad 10.32.220.19      0.0.0.0          "ras019"
20: fffc14 10:00:00:05:1e:34:00:63 10.32.220.20      0.0.0.0          >"ras020"
30: fffc1e 10:00:00:60:69:90:02:21 10.32.220.30      0.0.0.0          "ras030"
31: fffc1f 10:00:00:60:69:90:02:60 10.32.220.31      0.0.0.0          "ras031"
32: fffc20 10:00:00:60:69:90:02:68 10.32.220.32      0.0.0.0          "ras032"
33: fffc21 10:00:00:60:69:90:03:20 10.32.220.33      0.0.0.0          "ras033"
34: fffc22 10:00:00:60:69:90:03:01 10.32.220.34      0.0.0.0          "ras034"
40: fffc28 10:00:00:60:69:50:06:7f 10.32.220.40      0.0.0.0          "ras040"
45: fffc2d 10:00:00:05:1e:34:c5:17 10.32.220.45      0.0.0.0          "ras045"
50: fffc32 10:00:00:60:69:c0:06:64 10.32.220.50      0.0.0.0          "ras050"
51: fffc33 10:00:00:60:69:c0:1e:43 10.32.220.51      0.0.0.0          "ras051"
60: fffc3c 10:00:00:60:69:12:34:44 10.32.220.60      0.0.0.0          "ras060"
62: fffc3e 10:00:00:60:69:12:32:76 10.32.220.62      0.0.0.0          "ras062"
63: fffc3f 10:00:00:60:69:12:45:6e 10.32.220.63      0.0.0.0          "ras063"
64: fffc40 10:00:00:60:69:12:1d:51 10.32.220.64      0.0.0.0          "ras064"

The Fabric has 25 switches

ras045:admin>

```

The fields in the **fabricShow** display are:

Switch ID	The switch Domain_ID and embedded port D_ID.
Worldwide Name	The switch WWN.
Enet IP Addr	The switch Ethernet IP address.
FC IP Addr	The switch FC IP address.
Name	The switch symbolic name. An arrow (>) indicates the principal switch.

To set the domain ID

1. Connect to the switch and log in as admin.
2. Enter the **switchDisable** command to disable the switch.
3. Enter the **configure** command.
4. Enter **y** after the Fabric Parameters prompt:

```
Fabric parameters (yes, y, no, n): [no] y
```

5. Enter a unique domain ID at the Domain prompt. Use a domain ID value from 1 through 239 for normal operating mode (FCSW compatible):

```
Domain: (1..239) [1] 3
```

6. Respond to the remaining prompts (or press **Ctrl-d** to accept the other settings and exit).
7. Enter the **switchEnable** command to reenable the switch.

Linking Through a Gateway

A gateway merges SANs into a single fabric by establishing point-to-point E_Port connectivity between two Fibre Channel switches that are separated by a network with a protocol such as IP or SONET.

Except for link initialization, gateways are transparent to switches; the gateway simply provides E_Port connectivity from one switch to another.

By default, switch ports initialize links using the Exchange Link Parameters (ELP) mode 1. However, gateways expect initialization with ELP mode 2, also referred to as ISL R_RDY mode. Therefore, to enable two switches to link through a gateway, the ports on both switches must be set for ELP mode 2.

Any number of E_Ports in a fabric can be configured for gateway links, provided the following rules are followed:

- All switches in the fabric must be upgraded to Fabric OS v3.1.0 (or later) or v4.1.0 (or later).
- To prevent fabric segmentation, make sure that all switches in the fabric are using the core PID format, as described in [“To configure a link through a gateway”](#) next.
- When determining switch count maximums, include the switches connected to both sides of the gateway.
- Extended links (those created using the Extended Fabrics licensed feature) and the security features in Secure Fabric OS are not supported through gateway links.

To configure a link through a gateway

1. If you are not sure that the PID format is consistent across the entire fabric, enter the **configShow** command on all switches to check the PID setting. If necessary, change the PID format on any nonconforming switches, as described in [“Configuring the PID Format” on page A-1](#).
2. Connect to the switch on one end of the gateway and log in as admin.
3. Enter the **portCfgIslMode** command:

SilkWorm 200E, 3016, 3250, 3850, 3900, 4012, and 4100 switches:

portCfgIslMode *port mode*

Specify a *port* number. Valid values for port number vary depending on the switch type. The mode operand is required: specify 1 to enable ISL R_RDY mode (gateway link) or specify 0 to disable it.

SilkWorm 12000, 24000, and 48000 directors:

portCfgIslMode *slot/port, mode*

Specify a *slot/port* number pair. Valid values for slot and port number vary depending on the switch type. The mode operand is required: specify 1 to enable ISL R_RDY mode (gateway link) or specify 0 to disable it.

In this example, slot 2, port 3 is enabled for a gateway link:

```
switch:admin> portcfgislmode 2/3, 1
Committing configuration...done.
ISL R_RDY Mode is enabled for port 3. Please make sure the PID
formats are consistent across the entire fabric.
switch:admin>
```

4. Repeat the steps for any additional ports that will be connected to the gateway.

Repeat the procedure on the switch at the other end of the gateway.

Refer to the *Fabric OS Command Reference Manual* for more information about the **portCfgIslMode** command.

Checking Status

You can check the status of switch operation, high availability features, and fabric connectivity.

To check switch operation

1. Connect to the switch and log in as admin.
2. Enter the **switchShow** command at the command line. This command displays a switch summary and a port summary.
3. Check that the switch and ports are online.
4. Use the **switchStatusShow** command to further check the status of the switch.

To verify high-availability features

High-Availability (HA) features provide maximum reliability and nondisruptive replacement of key hardware and software modules. To verify these features, connect to the switch as admin and use any of the following commands:

- **chassisshow** verifies the field replaceable units (FRUs).
- **SilkWorm 12000, 24000, and 48000 directors**
 - **haShow** verifies that HA is enabled, that the heartbeat is up, and that the HA state is synchronized between the active and standby CP blades.
 - **slotShow** inventories and displays the current status of each slot in the system.

To verify fabric connectivity

1. Connect to the switch and log in as admin.
2. Enter the **fabricShow** command at the command line. This command displays a summary of all the switches in the fabric.

Example

```
switch:admin> fabricshow
Switch ID      Worldwide Name      Enet IP Addr      FC IP Addr      Name
-----
1: fffc01 10:00:00:60:69:80:04:5a 192.168.186.61 192.168.68.193 "switch61"
3: fffc03 10:00:00:60:69:10:9c:29 192.168.186.175 0.0.0.0 "switch175"
4: fffc04 10:00:00:60:69:12:14:b7 192.168.174.70 0.0.0.0 "switch70"
5: fffc05 10:00:00:60:69:45:68:04 192.168.144.121 0.0.0.0 "switch121"
6: fffc06 10:00:00:60:69:00:54:ea 192.168.174.79 192.168.68.197 "switch79"
7: fffc07 10:00:00:60:69:80:04:5b 192.168.186.62 192.168.68.194 "switch62"
8: fffc08 10:00:00:60:69:04:11:22 192.168.186.195 0.0.0.0 "switch195"
9: fffc09 10:00:00:60:69:10:92:04 192.168.189.197 192.168.68.198 "switch197"
10: fffc0a 10:00:00:60:69:50:05:47 192.168.189.181 192.168.68.181 "switch181"
11: fffc0b 10:00:00:60:69:00:54:e9 192.168.174.78 192.168.68.196 "switch78"
15: fffc0f 10:00:00:60:69:30:1e:16 192.168.174.73 0.0.0.0 "switch73"
33: fffc21 10:00:00:60:69:90:02:5e 192.168.144.120 0.0.0.0 "switch120"
44: fffc2c 10:00:00:60:69:c0:06:8d 192.168.144.121 0.0.0.0 "switch121"
97: fffc61 10:00:00:60:69:90:02:ed 192.168.144.123 0.0.0.0 "switch123"
98: fffc62 10:00:00:60:69:90:03:32 192.168.144.122 0.0.0.0 "switch122"

The Fabric has 15 switches

switch:admin>
```

To verify device connectivity

1. Connect to the switch and log in as admin.
2. *Optional:* Enter the **switchShow** command to verify that devices, hosts, and storage are connected.
3. *Optional:* Enter the **nsShow** command to verify that devices, hosts, and storage have successfully registered with the Name Server.
4. Enter the **nsAllShow** command at the command line. This command displays 24-bit Fibre Channel addresses of all devices in the fabric.

Example

```
switch:admin> nsallshow
{
  010e00 012fe8 012fef 030500 030b04 030b08 030b17 030b18
  030b1e 030b1f 040000 050000 050200 050700 050800 050de8
  050def 051700 061c00 071a00 073c00 090d00 0a0200 0a07ca
  0a07cb 0a07cc 0a07cd 0a07ce 0a07d1 0a07d2 0a07d3 0a07d4
  0a07d5 0a07d6 0a07d9 0a07da 0a07dc 0a07e0 0a07e1 0a0f01
  0a0f02 0a0f0f 0a0f10 0a0f1b 0a0f1d 0b2700 0b2e00 0b2fe8
  0b2fef 0f0000 0f0226 0f0233 0f02e4 0f02e8 0f02ef 210e00
  211700 211fe8 211fef 2c0000 2c0300 611000 6114e8 6114ef
  611600 620800 621026 621036 6210e4 6210e8 6210ef 621400
  621500 621700 621a00
  75 Nx_Ports in the Fabric }
switch:admin>
```

The number of devices listed should reflect the number of devices that are connected.

Tracking and Controlling Switch Changes

The track changes feature allows you to keep record of specific changes that might not be considered switch events, but might provide useful information. The output from the track changes feature is dumped to the system messages log for the switch. Use the **errDump** or **errShow** command to view the log.

Items in the log created from the Track changes feature are labeled TRACK.

Trackable changes are:

- Successful login
- Unsuccessful login
- Logout
- Configuration file change from task
- Track-changes on
- Track-changes off

An SNMP-TRAP mode can also be enabled; refer to the **trackChangesHelp** command in the *Fabric OS Command Reference Manual*.

For troubleshooting information on the track changes feature, refer to [“Inaccurate Information in the System Message Log” on page 10-23](#).

To enable the track changes feature

1. Connect to the switch and log in as admin.
2. Enter this command to enable the track changes feature: **trackChangesSet 1**.

A message displays, verifying that the track changes feature is on:

```
switch:admin> trackchangesset 1
Committing configuration...done.
switch:admin>
```

The output from the track changes feature is dumped to the system message log for the switch. Use the **errDump** or **errShow** command to view the log.

Items in the system message log created from the track changes feature are labeled TRCK:

```
2004/08/24-08:45:43, [TRCK-1001], 212,, INFO, ras007, Successful login by user
admin.
```

To display the status of the track changes feature

1. Connect to the switch and log in as admin.
2. Enter the **trackChangesShow** command.

The status of the track changes feature is displayed as either on or off. The display includes whether the track changes feature is configured to send SNMP traps:

```
switch:admin> trackchangesshow
Track changes status: ON
Track changes generate SNMP-TRAP: NO
switch:admin>
```

To view the switch status policy threshold values

1. Connect to the switch and log in as admin.
2. Enter the **switchStatusPolicyShow** command at the command line.

Whenever there is a switch change, an error message is logged and an SNMP connUnitStatusChange trap is sent.

SilkWorm 200E, 3250, 3850, 3900, 4012, and 4100 switches: The output is similar to the following:

```
switch:admin> switchstatuspolicyshow
The current overall switch status policy parameters:
              Down      Marginal
-----
PowerSupplies 2          1
Temperatures  2          1
Fans          2          1
Flash         0          1
MarginalPorts 5          2
FaultyPorts   2          1
MissingSFPS   2          1
switch:admin>
```

SilkWorm 3016: Because this model has no power supplies or fans, they are not listed in the output.

SilkWorm 12000, 24000 and 48000: The output is similar to the following:

```
switch:admin> switchstatuspolicyshow
The current overall switch status policy parameters:
      Down      Marginal
-----
PowerSupplies  3          0
Temperatures   2          1
  Fans         2          1
    WWN        0          1
    CP         0          1
    Blade      0          1
    Flash      0          1
MarginalPorts  2          1
  FaultyPorts  2          1
  MissingSFPs  0          0
switch:admin>
```

The policy parameter determines the number of failed or inoperable units for each contributor that will trigger a status change in the switch.

Each parameter can be adjusted so that a specific threshold must be reached before that parameter changes the overall status of a switch to MARGINAL or DOWN. For example, if the FaultyPorts DOWN parameter is set to 3, the status of the switch will change if 3 ports fail. Only one policy parameter needs to pass the MARGINAL or DOWN threshold to change the overall status of the switch.

For more information about setting policy parameters, refer to the *Fabric Watch Administrator's Guide*.

To set the switch status policy threshold values

1. Connect to the switch and log in as admin.
2. Enter the **switchStatusPolicySet** command at the command line.

First, the current switch status policy parameter values are displayed. Then, you are prompted to enter values for each DOWN and MARGINAL threshold parameter:

3. Verify the threshold settings you have configured for each parameter. Enter the **switchStatusPolicyShow** command to view your current switch status policy configuration:



Note

By setting the DOWN and MARGINAL value for a parameter to 0,0 that parameter is no longer used in setting the overall status for the switch.

SilkWorm 3016, 3250, 3850, 3900, 4012, and 4100 switches: The following example shows the command as executed on a SilkWorm 3900 switch. The output is similar on SilkWorm 3016, 3250, 3850, 4012, and 4100 models:

```
switch:admin> switchstatuspolicyset
To change the overall switch status policy parameters
The current overall switch status policy parameters:
      Down      Marginal
-----
FaultyPorts  2      1
MissingSFPS  0      0
PowerSupplies 2      1
Temperatures 2      1
Fans         2      1
PortStatus   0      0
ISLStatus    0      0
Note that the value, 0, for a parameter, means that it is
NOT used in the calculation.
** In addition, if the range of settable values in the prompt is (0..0),
** the policy parameter is NOT applicable to the switch.
** Simply hit the Return key.
The minimum number of
  FaultyPorts contributing to
                        DOWN status: (0..32) [2] 3
  FaultyPorts contributing to
                        MARGINAL status: (0..32) [1] 2
  MissingSFPS contributing to
                        DOWN status: (0..32) [0]
  MissingSFPS contributing to
                        MARGINAL status: (0..32) [0]
  Bad PowerSupplies contributing to
                        DOWN status: (0..2) [2]
  Bad PowerSupplies contributing to
                        MARGINAL status: (0..2) [1]
  Bad Temperatures contributing to
                        DOWN status: (0..5) [2]
  Bad Temperatures contributing to
                        MARGINAL status: (0..5) [1]
  Bad Fans contributing to
                        DOWN status: (0..6) [2]
  Bad Fans contributing to
                        MARGINAL status: (0..6) [1]
  Down PortStatus contributing to
                        DOWN status: (0..32) [0]
  Down PortStatus contributing to
                        MARGINAL status: (0..32) [0]
  down ISLStatus contributing to
                        DOWN status: (0..32) [0]
  down ISLStatus contributing to
                        MARGINAL status: (0..32) [0]
Policy parameter set has been changed
```

SilkWorm 3016: Command output does not include power supplies or fan information.

SilkWorm 12000, 24000, and 48000: Command output includes parameters related to CP blades.

Configuring Standard Security Features

This chapter provides information and procedures for configuring standard Fabric OS security features such as account and password management.

Additional security features are available when secure mode is enabled. For information about licensed security features available in Secure Fabric OS, refer to the *Secure Fabric OS Administrator's Guide*.

Secure Protocols

Fabric OS supports the secure protocols shown in [Table 3-1](#).

Table 3-1 Secure Protocol Support

Protocol	Description
SSL	Supports SSLv3, 128-bit encryption by default. Fabric OS uses SSL to support HTTPS. A certificate must be generated and installed on each switch to enable SSL.
HTTPS	Web Tools supports the use of HTTPS.
Secure File Copy (scp)	Configuration upload and download support the use of scp.
SNMPv3	SNMPv1 is also supported.

Simple Network Management Protocol (SNMP) is a standard method for monitoring and managing network devices. Using SNMP components, you can program tools to view, browse, and manipulate Brocade switch variables and set up enterprise-level management processes.

Every Brocade switch carries an SNMP agent and management information base (MIB). The agent accesses MIB information about a device and makes it available to a network manager station. You can manipulate information of your choice by *trapping* MIB elements using the Fabric OS CLI, Web Tools, or Fabric Manager.

The SNMP Access Control List (ACL) provides a way for the administrator to restrict SNMP get/set operations to certain hosts/IP addresses. This is used for enhanced management security in the storage area network.

For details on Brocade MIB files, naming conventions, loading instructions, and information about using Brocade's SNMP agent, refer to the *Fabric OS MIB Reference Manual*.

[Table 3-2](#) describes additional software or certificates that you must obtain to deploy secure protocols.

Table 3-2 Items Needed to Deploy Secure Protocols

Protocol	Host Side	Switch Side
Secure telnet (sectelnet)	Sectelnet client	License not required, but a switch certificate issued by Brocade is required
SSH	SSH client	None
HTTPS	No requirement on host side except a browser that supports HTTPS	Switch IP certificate for SSL
Secure File Copy (scp)	SSH daemon, scp server	None
SNMPv3, SNMPv1	None	None

The security protocols are designed with the four main usage cases described in [Table 3-3](#).

Table 3-3 Main Security Scenarios

Fabric	Management Interfaces	Comments
Nonsecure	Nonsecure	No special setup is need to use telnet or HTTP. A Brocade switch certificate must be installed if sectelnet is used.
Nonsecure	Secure	Secure protocols may be used. An SSL switch certificate must be installed if SSH/HTTPS is used.
Secure	Secure	Secure protocols are supported on Fabric OS v4.4.0 (and later) switches. Switches running earlier Fabric OS versions can be part of the secure fabric, but they do not support secure management. Secure management protocols must be configured for each participating switch. Nonsecure protocols may be disabled on nonparticipating switches. If SSL is used, then certificates must be installed.
Secure	Nonsecure	You must use sectelnet because telnet is not allowed in secure mode. Nonsecure management protocols are necessary under these circumstances: <ul style="list-style-type: none"> The fabric contains switches running Fabric OS v3.2.0. The presence of software tools that do not support Secure protocols: for example, Fabric Manager v4.0.0. The fabric contains switches running Fabric OS versions earlier than v4.4.0. Nonsecure management is enabled by default.

Ensuring Network Security

To ensure security, Fabric OS supports secure shell (SSH) encrypted sessions. SSH encrypts all messages, including the client's transmission of password during login. The SSH package contains a daemon (sshd), which runs on the switch. The daemon supports a wide variety of encryption algorithms, such as Blowfish-CBC and AES.



Note

To maintain a secure network, you should avoid using telnet or any other unprotected application when you are working on the switch. For example, if you use telnet to connect to a machine, and then start an SSH or secure telnet session from that machine to the switch, the communication to the switch is in clear text and therefore is not secure.

The FTP protocol is also not secure. When you use FTP to copy files to or from the switch, the contents are in clear text. This includes the remote FTP server's login and password. This limitation affects the following commands: **saveCore**, **configUpload**, **configDownload**, and **firmwareDownload**.

Commands that require a secure login channel must be issued from an original SSH session. If you start an SSH session, and then use the login command to start a nested SSH session, commands that require a secure channel will be rejected.

Fabric OS v4.4.0 and later supports SSH protocol v2.0 (ssh2). For more information on SSH, refer to the SSH IETF Web site:

<http://www.ietf.org/ids.by.wg/secsh.html>

Refer to *SSH, The Secure Shell: The Definitive Guide* by Daniel J. Barrett, Richard Silverman.

Fabric OS v4.4.0 comes with the SSH server preinstalled; however, you must select and install the SSH client. For information on installing and configuring the F-Secure SSH client, refer to the Web site:

<http://www.f-secure.com>

Configuring the Telnet Interface

Telnet is enabled by default. To prevent users from passing clear text passwords over the network when they connect to the switch, you can disable the telnet interface.



Note

Before disabling the telnet interface, make sure you have an alternate method of establishing a connection with the switch.

To disable telnet

1. Connect to the switch and log in as admin.

Connect through some other means than telnet: for example, through SSH.

2. Enter the following command:

```
configure telnetd
```

3. In response to the System Services prompt, enter **y**.
4. In response to the telnetd prompt, enter **off**.

The telnet interface is disabled. If you entered the command during a standard telnet session, the session terminates.

Example

```
switch:admin> configure telnetd
Not all options will be available on an enabled switch.
To disable the switch, use the "switchDisable" command.
Configure...
ssl attributes (yes, y, no, n): [no]
http attributes (yes, y, no, n): [no]
snmp attributes (yes, y, no, n): [no]
rpcd attributes (yes, y, no, n): [no]
cfgload attributes (yes, y, no, n): [no]

[31454]: Read 1 license entries for generation 1.
[31454]: Read 1 license records.
System services (yes, y, no, n): [no] y

rstatd (on, off): [off]
rusersd (on, off): [off]
telnetd (on, off): [on] off
```

To enable telnet

1. Connect to the switch through a means other than telnet (for example, SSH) and log in as admin.
2. Enter the following command:

```
configure telnetd
```

3. In response to the System Services prompt, enter **y**.
4. In response to the telnetd prompt, enter **on**.

The telnet interface is enabled.

Blocking Listeners

SilkWorm switches block Linux subsystem listener applications that are not used to implement supported features and capabilities. [Table 3-4](#) lists the listener applications that SilkWorm switches either block or do not start.

Table 3-4 Blocked Listener Applications

Listener Application	SilkWorm 12000, 24000, and 48000 Directors	SilkWorm 200E, 3016, 3250, 3850, 3900, 4012, and 4100 Switches
chargen	Do not start	Do not start
echo	Do not start	Do not start
daytime	Do not start	Do not start
discard	Do not start	Do not start
ftp	Do not start	Do not start
rexec	Block with packet filter	Do not start
rsh	Block with packet filter	Do not start
rlogin	Block with packet filter	Do not start
time	Block with packet filter	Do not start
rstats	Do not start	Do not start
rusers	Do not start	Do not start

Accessing Switches and Fabrics

[Table 3-5](#) lists the defaults for accessing hosts, devices, switches, and zones.

Table 3-5 Access Defaults

Hosts	Any host can access the fabric by SNMP
	Any host can telnet to any switch in the fabric
	Any host can establish an HTTP connection to any switch in the fabric
	Any host can establish an API connection to any switch in the fabric
Devices	All device ports can access SES
	All devices can access the management server
	Any device can connect to any FC port in the fabric
Switch access	Any switch can join the fabric
	All switches in the fabric can be accessed through serial port
Zoning	Node WWNs can be used for WWN-based zoning

Creating and Maintaining User-Defined Accounts

In addition to the default administrative and user accounts, Fabric OS supports up to 15 user-defined accounts in each logical switch (domain). These accounts expand your ability to track account access and audit administrative activities.

User-defined accounts can be assigned either admin, switchAdmin, or user level roles. Admin-level accounts allow up to two simultaneous login sessions. User-level accounts allow up to four simultaneous login sessions. The total number of simultaneous login sessions allowed per logical switch is 15.

You can change passwords on user-defined accounts as described in [“Changing an Account Password” on page 3-8](#).

If the track changes feature is enabled, the system keeps track of account names and login attempts. (Refer to [“Tracking and Controlling Switch Changes” on page 2-20](#) for details on enabling the track changes feature.)

For large enterprises, Fabric OS also supports RADIUS services, as described in [“Setting Up RADIUS AAA Service” on page 3-9](#).

The following procedures are for operations you can perform on user-defined accounts.

SilkWorm 3016: The default administrative account is called “USERID”. For instructions on changing this name, refer to [“To change the default administrative account name on the SilkWorm 3016” on page 2-5](#). On all other models, the default administrative account is “admin”.



Note

If you are operating in secure mode, you can perform these operations only on the primary FCS switch.

To display account information

1. Connect to the switch and log in as admin.
2. Enter one of the following commands:
 - `userConfig --show -a` to show all account information for a logical switch
 - `userConfig --show -b` to show all backup account information for a logical switch
 - `userConfig --show username` to show account information for the specified account name

Accounts with the admin role can display information about all accounts on the logical switch.

Accounts with the switchAdmin role can display information about all accounts on the logical switch, however, it cannot display information about security, user management, or zoning. Accounts with the user role can display only information about themselves.

To create a user-defined account

1. Connect to the switch and log in as admin.
2. Enter the following command:

```
userConfig --add username -r rolename [-d description]
```

<i>username</i>	Specifies the account name, which must begin with an alphabetic character. The name can be from 8 to 40 characters long. It is case-sensitive and can contain alphabetic and numeric characters, the dot (.) and the underscore (_). It must be different than all other account names on the logical switch.
<i>-r rolename</i>	Specifies the role: either admin, switchAdmin, or user in nonsecure mode; admin, user, or nonfcsadmin in secure mode.
<i>-d description</i>	Optionally, adds a description to the account. The description field can be up to 40 printable ASCII characters long. The following characters are not allowed: asterisk (*), quotation mark ("), exclamation point (!), semicolon (;), and colon (:).

3. In response to the prompt, enter a password for the account.

The password is not displayed when you enter it on the command line.

Accounts with the admin role can create accounts. Accounts with the user role cannot.

To delete a user-defined account

1. Connect to the switch and log in as admin.
2. Enter the following command:

```
userConfig --delete username
```

<i>username</i>	Specifies the account name. You cannot delete the default accounts. An account cannot delete itself. All active CLI sessions for the deleted account are logged out.
-----------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------

3. At the prompt for confirmation, enter **y**.

Only accounts with the admin role can delete user-defined accounts on the logical switch.

To change account parameters

1. Connect to the switch and log in as admin.
2. Enter the following command:

```
userconfig --change username [-r rolename] [-d description] [-e yes | no]
```

<i>username</i>	Changes the account attribute for username. The account must already exist.
<i>-r rolename</i>	Optionally, changes the role: either admin, switchAdmin, or user in nonsecure mode; admin, user, or nonfcsadmin in secure mode. An account cannot change its own role. You can change only the role name of a user-defined account with a lower level of authorization.
<i>-d description</i>	Optionally, the account description. The description field can be up to 40 printable ASCII characters long. The following characters are not allowed: asterisk (*), quotation mark ("), exclamation point (!), semicolon (;), and colon (:). You can change only the description of a user-defined account with a lower level of authorization.
<i>-e</i>	Optionally, enter yes to enable the account or enter no to disable it. If you disable an account, all active CLI sessions for that account are logged out. You can enable or disable user-defined or default accounts.

Accounts with the admin role can change information for accounts that have lesser permissions. Accounts with the user role cannot.

To recover user-defined accounts

If a backup account exists (in secure mode), you can recover it with the following command:

```
userConfig --recover
```

The following conditions apply to recovering user accounts:

- Only accounts with admin or higher roles can recover accounts.
- The attributes in the backup database replace the attributes in the current account database.
- An event is stored in the system message log, indicating that accounts have been recovered.

Changing an Account Password

At each level of account access, you can change passwords for that account and accounts that have lesser privileges.

If you log in to a user account, you can change only that account's password.

If you log in to an admin account, you can change admin and user passwords. You must provide the old password when the account being changed has the same or higher privileges than the current login account. For example, when logged in as admin, you need admin passwords to change passwords for admin accounts (except when you change the default user account password at login), but you do not need user passwords to change passwords for user accounts.

A new password must have at least one character different from the old password. The following rules also apply to passwords:

- You cannot change passwords using SNMP.
- Password prompting is disabled when security mode is enabled.
- Starting with Fabric OS v4.4.0, admin-level accounts can use Web Tools to change passwords.
- Starting with Fabric OS v3.2.0, you cannot change default account names.

For information on password behavior when you upgrade (or downgrade) firmware, refer to [“Effects of Firmware Changes on Accounts and Passwords” on page 4-7](#).

To change the password for the current login account

1. Connect to the switch and log in as either admin or user.
2. Enter the following command:

```
passwd
```

Enter the requested information at the prompts.

To change the password for a different account

1. Connect to the switch and log in as admin.
2. Enter the following command:

```
passwd name
```

where *name* is the name of the account.

Enter the requested information at the prompts.

If the named account has lesser privileges than the current login account, the old password of the named account is not required. If the named account has equal or higher privileges than the current login account, you are prompted to enter the old password of the named account.

Setting Up RADIUS AAA Service

Fabric OS v3.2, v4.4.0, and later support RADIUS authentication, authorization, and accounting service (AAA). When configured for RADIUS, a switch becomes a RADIUS client. In this configuration, authentication records are stored in the RADIUS host server database. Login and logout account name, assigned role, and time-accounting records are also stored on the RADIUS server for each user.

By default, RADIUS service is disabled, so AAA services default to the switch local database.

To enable RADIUS service, it is strongly recommended that you access the CLI through an SSH connection so that the shared secret is protected. Multiple login sessions can configure simultaneously, and the last session to apply a change leaves its configuration in effect. After a configuration is applied, it persists after a reboot or an HA failover.

The configuration is chassis-based, so it applies to all logical switches (domains) on the switch and replicates itself on a standby CP blade, if one is present. It is saved in a configuration upload and applied in a configuration download.

You should configure at least two RADIUS servers so that if one fails, the other will assume service. You can set the configuration with both RADIUS service and local authentication enabled so that if all RADIUS servers do not respond (because of power failure or network problems), the switch uses local authentication.

Consider the following effects of the use of RADIUS service on other Fabric OS features:

- When RADIUS service is enabled, all account passwords must be managed on the RADIUS server. The Fabric OS mechanisms for changing switch passwords remain functional; however, such changes affect only the involved switches locally. They do not propagate to the RADIUS server, nor do they affect any account on the RADIUS server.

When RADIUS is set up for a fabric that contains a mix of switches with and without RADIUS support, the way a switch authenticates users depends on whether or not a RADIUS server is set up for that switch. For a switch with RADIUS support and configuration, authentication bypasses the local password database. For a switch without RADIUS support or configuration, authentication uses the switch's local account names and passwords.

- When Secure Fabric OS secure mode is enabled, the following behaviors apply:
 - Account passwords stored in the switch-local password database are distributed among all switches in the same fabric. RADIUS configuration is not affected.
 - There are separate admin and nonfcsadmin roles in secure mode. A nonfcsadmin account on a RADIUS server cannot access FCS switches, even if the account is properly authenticated.
 - If a nonfcsadmin account on a RADIUS server logs in to a switch in nonsecure mode, the switch grants the user admin role privileges.
- The following behaviors apply to Web Tools:
 - Web Tools client and server keep a session open after a user is authenticated. A password change on a switch invalidates an open session and requires the user to log in again. When integrated with RADIUS, a switch password change on the RADIUS server does not invalidate an existing open session, although a password change on the local switch does.
 - If you cannot log in because of a RADIUS server connection problem, Web Tools displays a message indicating server outage.

Configuring the RADIUS Server

You must know the switch IP address or name to connect to switches. Use the **ipAddrShow** command to display a switch IP address.

For SilkWorm directors (chassis-based systems), the switch IP addresses are aliases of the physical Ethernet interfaces on the CP blades. When specifying client IP addresses for the logical switches in such systems, make sure that the CP blade IP addresses are used. For accessing both the active and standby CP blade, and for the purpose of HA failover, both of the CP blade IP addresses should be included in the RADIUS server configuration.

User accounts should be set up by their true network-wide identity, rather than by the account names created on a Fabric OS switch. Along with each account name, the administrator should assign appropriate switch access roles. To manage a nonsecure fabric, these roles can be user or admin. To manage a secure fabric, these roles can be user, admin, or nonfcsadmin.

When they log in to a switch configured with RADIUS, users enter their assigned RADIUS account names and passwords at the prompt. After RADIUS server authenticates a user, it responds with the assigned switch role in a Brocade Vendor-Specific Attribute (VSA), as defined in the RFC. An Authentication-Accept response without such VSA role assignment automatically assigns the user role.

The following sections describe how to configure a RADIUS server to support Brocade clients under different operating systems.

Linux

The following procedures work for FreeRADIUS on Solaris and Red Hat Linux. FreeRADIUS is a freeware RADIUS server that you can find at:

www.freeradius.org

Follow the installation instructions at the Web site. FreeRADIUS runs on Linux (all versions), FreeBSD, NetBSD, and Solaris. If you make a change to any of the files used in this configuration, you must stop the server and restart it for the changes to take effect.

FreeRADIUS installation places the configuration files in *\$PREFIX/etc/raddb*. By default, the PREFIX is */usr/local*.

Configuring RADIUS service on Linux consists of the following tasks:

- Adding the Brocade attribute to the server
- Creating the user
- Enabling clients

To add the Brocade attribute to the server

1. Create and save the file *\$PREFIX/etc/raddb/dictionary.brocade* with the following information:

```
#
# Brocade FabricOS v5.0.1 dictionary
#
VENDOR      Brocade          1588
#
# attribute 1 defined to be Brocade-Auth-Role
# string defined in user configuration
#
ATTRIBUTE Brocade-Auth-Role 1      string      Brocade
```

This defines the Brocade vendor ID as 1588, the Brocade attribute 1 as Brocade-Auth-Role, and it is a string value.

2. Open the file *\$PREFIX/etc/raddb/dictionary* in a text editor and add the line:

```
$INCLUDE dictionary.brocade
```

As a result, the file *dictionary.brocade* is located in the RADIUS configuration directory., and loaded for use by the RADIUS server.

To create the user

Open the `$PREFIX/etc/raddb/user` file in a text editor and add user names and roles for users who will be accessing the switch and authenticating RADIUS. The user will log in using the role specified with `Brocade-Auth-Role`. The valid roles include `root`, `factory`, `admin`, `switchAdmin`, and `user`. You must use quotation marks around “password” and “role”.

For example, to set up an account called `JohnDoe` with the `admin` role:

```
JohnDoe Auth-Type := Local, User-Password == "johnPassword" Brocade-Auth-Role =
"admin"
```

The next example uses the local system password file to authenticate users. (This does not work when using NIS for authentication. The only way to enable authentication with the password file is to force the Brocade switch to authenticate using PAP; this requires the `-a pap` option with the **aaaConfig** command.)

```
JohnDoe Auth-Type := System, Brocade-Auth-Role = "admin"
```

To enable clients

Clients are the switches that will be using the RADIUS server; each client must be defined. By default, all IP addresses are blocked.

On dual-CP switches (SilkWorm 12000, 24000, and 48000), the switch sends its RADIUS request using the IP address of the active CP. When adding clients, add both the active and standby CP IP addresses so that users can still log in in the event of a failover.

1. Open the `$PREFIX/etc/raddb/client.config` file in a text editor and add the switches that are to be configured as RADIUS clients. For example, to configure the switch at IP address 10.32.170.59 as a client:

```
client 10.32.170.59
  secret      = Secret
  shortname   = Testing Switch
  nastype     = other
```

In this example, *shortname* is an alias used to easily identify the client. *Secret* is the shared secret between the client and server. Make sure that the shared secret matches that configured on the switch (see [“To add a RADIUS server to the switch configuration” on page 3-16](#)).

2. Save the file `$PREFIX/etc/raddb/client.config`; then start the RADIUS server as follows:

```
$PREFIX/sbin/radiusd
```

Windows 2000

Configuring RADIUS service on Windows 2000 consists of the following tasks:

- Installing Internet Authentication Service (IAS)

For more information and instructions on installing IAS, refer to the Microsoft Web site.

- Enabling the Challenge Handshake Authentication Protocol (CHAP)

If CHAP authentication is required, then Windows must be configured to store passwords with reversible encryption. Reverse password encryption is not the default behavior; it must be enabled.



Note

If a user is configured prior to enabling reverse password encryption, then the user's password is stored and cannot utilize CHAP. To use CHAP, the password must be reentered after encryption is enabled. If the password is not reentered, then CHAP authentication will not work and the user will be unable to authenticate from the switch.

- Configuring a user

Internet Authentication Service (IAS) is the Microsoft implementation of a RADIUS server and proxy. IAS uses the Windows native user database to verify user login credentials; it does not list specific users, but instead lists *user groups*. Each user group should be associated with a specific switch login role. For example, you should configure a user group for root, admin, factory, switchAdmin, and user, and then add any users whose logins you want to associate to the appropriate group.

- Configuring the server

To enable CHAP

1. From the Windows Start menu, select **Programs> Administrative Tools> Local Security Policy** to open the Local Security Settings window.
2. In the Local Security Settings window, expand the **Account Policies** folder and select the **Password Policy** folder. From the list of policies in the Password Policy folder, right-click **Store password using reversible encryption for all users in the domain**, and select **Security** from the pop-up menu.
3. An additional Local Security Settings window appears. Click the **Enabled** radio button and then click **OK**.

To configure users

1. From the Windows Start menu, select **Programs> Administrative Tools>Computer Management** to open the Computer Management window.
2. In the Computer Management window, expand the **Local Users and Groups** folder and select the **Groups** folder. Right-click the **Groups** folder and select **New Group** from the pop-up menu.
3. In the New Group window, provide a Name and Description for the group and click **Add**.
4. In the Select Users or Groups window, select the user—who should already have been configured—you want to add to the group and click **Add**. Repeat this for every user you want to add. When you have completed adding all users, click **OK**.
5. In the New Group window, verify that the users you added in step 4 appear in the Members field; then click **Create** to create this group.

The new groups are created for each login type (admin, switchAdmin, user).

To configure the RADIUS server

1. From the Windows Start menu, select **Programs> Administrative Tools>Internet Authentication Service** to open the Internet Authentication Service window.
2. In the Internet Authentication Service window, right-click the **Clients** folder and select **New Client** from the pop-up menu.



Note

A *client* is the device that uses the RADIUS server; in this case, it is the switch.

3. In the Add Client window, provide the following:

Friendly name	The Friendly name should be an alias that is easily recognizable as the switch to which you are connecting.
Protocol	Select RADIUS as the protocol.
4. In the Add RADIUS Client window, provide the following:

Client address (IP or DNS)	Enter the IP address of the switch.
Client-Vendor	Select RADIUS Standard .
Shared secret	Provide a password. Shared secret is a password used between the client device and server to prevent IP address spoofing by unwanted clients. Keep your shared secret password in a safe place. You will need to enter this password in the switch configuration.

After clicking **Finish**, repeat steps 2 through 4 for all switches on which RADIUS authentication will be used.
5. In the Internet Authentication Service window, right-click the Remote Access Policies folder; then select **New Remote Access Policy** from the pop-up window.

A remote access policy must be created for each Brocade login role (root, admin, factory, switchAdmin, and user) for which you want to use RADIUS. Apply this policy to the user groups that you already created.
6. In the Add Remote Access Policy window, enter an easily identifiable **Policy friendly name** that will enable you to see the switch login for which the policy is being created; then click **Next**.
7. After the Add Remote Access Policy window refreshes, click **Add**.
8. In the Select Attribute window, select **Windows Groups** and click **Add**.
9. In the Groups window, click **Add**.
10. In the Select Groups window, select the user-defined group for which you are creating a policy and click **Add**. After adding all appropriate groups, click **OK**. In the Groups window, click **OK**.
11. In the Add Remote Access Policy window, confirm that the “Conditions” section displays the group(s) that you selected and click **Next**.

After the Add Remote Access Policy window refreshes, select the **Grant remote access permission** radio button and click **Next**.

After the Add Remote Access Policy window refreshes again, click **Edit Profile**.

12. In the Edit Dial-in Profile window, click the **Authentication** tab and check only the **Encrypted Authentication (CHAP)** and **Unencrypted Authentication (PAP, SPAP)** checkboxes; then click the **Advanced** tab and click **Add**.
13. In the Add Attributes window, select **Vendor-Specific** and click **Add**.
14. In the Multivalued Attribute Information window, click **Add**.
15. In the Vendor-Specific Attribute Information window, click the **Enter Vendor Code** radio button and enter the value **1588**. Click the **Yes. It conforms** radio button, and then click **Configure Attribute....**
16. In the Configure VSA (RFC compliant) window, enter the following:

Vendor-assigned attribute number Enter the value **1**.

Attribute format Enter **String**.

Attribute value Enter the login role (root, admin, factory, switchAdmin, or user) the user group must use to log in to the switch.

Click **OK**.
17. In the Multivalued Attribute Information window, click **OK**.
18. In the Edit Dial-in Profile window, remove all additional parameters (except the one you just added, “Vendor-Specific”). Click **OK**.
19. In the Add Remote Access Policy window, click **Finish**. After returning to the Internet Authentication Service window, repeat steps 5 through 19 to add additional policies for all Brocade login types you want to use the RADIUS server. After this is done, you can configure the switch.

Configuring the Switch

RADIUS configuration of the switch is controlled by the **aaaConfig** command.



Note

On dual-CP switches (SilkWorm 12000, 24000, and 48000), the switch sends its RADIUS request using the IP address of the active CP. When adding clients, add both the active and standby CP IP addresses so that users can still log in in the event of a failover.

The following procedures show how to use the **aaaConfig** command to set up a switch for RADIUS service.

To display the current RADIUS configuration

1. Connect to the switch and log in as admin.
2. Enter this command:

```
switch:admin> aaaConfig --show
```

If a configuration exists, its parameters are displayed. If RADIUS service is not configured, only the parameter heading line is displayed. Parameters include:

Position	The order in which servers are contacted to provide service
Server	The server names or IP addresses
Port	The server ports
Secret	The shared secrets
Timeouts	The length of time servers have to respond before the next server is contacted
Authentication	The type of authentication being used on servers

To add a RADIUS server to the switch configuration

1. Connect to the switch and log in as admin.
2. Enter this command:

```
switch:admin> aaaConfig --add server [-p port] [-s secret] [-t timeout]
[-a pap | chap]
```

<i>server</i>	Enter either a server name or IP address. Avoid duplicating server listings (that is, listing the same server once by name and again by IP address). Up to five servers can be added to the configuration.
<i>-p port</i>	Optionally, enter a server port. The default is port 1812.
<i>-s secret</i>	Optionally, enter a shared secret. The default is “sharedsecret”. Secrets can be from 8 to 40 alphanumeric characters long. Make sure that the secret matches that configured on the server.
<i>-t timeout</i>	Optionally, enter the length of time (in seconds) that the server has to respond before the next server is contacted. The default is three seconds. Timeout values can range from 1 to 30 seconds.
<i>-a[pap chap]</i>	Specify PAP or CHAP as authentication protocol.

To enable or disable RADIUS service

1. Connect to the switch and log in as admin.
2. Enter this command:

```
switch:admin> aaaConfig --radius on | off
```

Specifying **on** enables the service; specifying **off** disables it.

At least one RADIUS server must be configured before you can enable RADIUS service.

If no RADIUS configuration exists, turning it on triggers an error message. When the command succeeds, the event log indicates that the configuration is enabled or disabled.

To delete a RADIUS server from the configuration

1. Connect to the switch and log in as admin.
2. Enter this command:

```
switch:admin> aaaConfig --remove server | all
```

server Servers are listed by either name or IP address. Enter either the name or IP address of the server to be removed.

3. At the prompt, enter **y** to complete the command.

When the command succeeds, the event log indicates that the server is removed.

To change a RADIUS server configuration

1. Connect to the switch and log in as admin.
2. Enter this command:

```
switch:admin> aaaConfig --change server [-p port] [-s secret] [-t timeout] [-a pap | chap]
```

server Servers are listed by either name or IP address. Enter either the name or IP address of the server to be changed.

-p port Optionally, enter a server port.

-s secret Optionally, enter a shared secret.

-t timeout Optionally, enter the length of time (in seconds) the server has to respond before the next server is contacted.

-a[pap|chap] Specify PAP or CHAP as authentication protocol.

To change the order in which RADIUS servers are contacted for service

1. Connect to the switch and log in as admin.
2. Enter this command:

```
switch:admin> aaaConfig --move server to_position
```

server Servers are listed by either name or IP address. Enter either the name or IP address of the server whose position is to be changed.

to_position Enter the position number to which the server is to be moved.

When the command succeeds, the event log indicates that a server configuration is changed.

Enabling and Disabling Local Authentication

It is useful to enable local authentication so that the switch can take over authentication locally if the RADIUS servers fail to respond because of power outage or network problems. To enable or disable local authentication, enter the following command:

```
switch:admin> aaaConfig --switchdb on | off
```

Specifying **on** enables local authentication; specifying **off** disables it.

When local authentication is enabled and RADIUS servers fail to respond, you can log in to the default switch accounts (admin and user) or any user-defined account. You must know the passwords of these accounts.

RADIUS authentication must be enabled when local database authentication is turned off from the on state; otherwise, an error is returned.

Because local database authentication might be automatically disabled or enabled when enabling or disabling RADIUS authentication, you should set the local database authentication explicitly to enabled or disabled *after* setting the desired RADIUS authentication configuration.

When the command succeeds, the event log indicates that local database authentication is disabled or enabled.

Configuring for the SSL Protocol

Fabric OS v4.4.0 and later supports secure sockets layer (SSL) protocol, which provides secure access to a fabric through Web-based management tools like Web Tools. SSL support is a standard Fabric OS feature; it is independent of Secure Fabric OS, which requires a license and separate certification.

Switches configured for SSL grant access to management tools through hypertext transfer protocol-secure links (which begin with *https://*) instead of standard links (which begin with *http://*).

SSL uses public key infrastructure (PKI) encryption to protect data transferred over SSL connections. PKI is based on digital certificates obtained from an Internet Certificate Authority (CA), which acts as the trusted key agent.

Certificates are based on the switch IP address or fully qualified domain name (FQDN), depending on the issuing CA. If you change a switch IP address or FQDN after activating an associated certificate, you might have to obtain and install a new certificate. Check with the CA to verify this possibility, and plan these types of changes accordingly.

Browser and Java Support

Fabric OS supports the following Web browsers for SSL connections:

- Internet Explorer (Microsoft Windows)
- Mozilla (Solaris and Red Hat Linux)

In countries that allow the use of 128-bit encryption, you should use the latest version of your browser. For example, Internet Explorer 6.0 and later supports 128-bit encryption by default. You can display the encryption support (called “cipher strength”) using the Internet Explorer **Help:About** menu option. If you are running an earlier version of Internet Explorer, you might be able to download an encryption patch from the Microsoft Web site at <http://www.microsoft.com>.

You should upgrade to the Java 1.4.2_03 Plug-in on your management workstation. To find the Java version that is currently running, open the Java console and look at the first line of the window.

For more details on levels of browser and Java support, refer to the *Web Tools Administrator's Guide*.

Summary of SSL Procedures

You configure for SSL by obtaining, installing, and activating digital certificates for SSL support. Certificates are required on all switches that are to be accessed through SSL.

You also need to install a certificate to the Java Plug-in on the management workstation, and you might need to add a certificate to your Web browser.

Configuring for SSL involves these major steps, which are shown in detail in the next sections:

1. Choose a CA.
2. On each switch:
 - a. Generate a public/private key (**secCertUtil genkey** command).
 - b. Generate a certificate signing request (CSR) (**secCertUtil genscr** command) and store the CSR on an FTP server (**secCertUtil export** command).
3. Obtain the certificates from the CA.

You can request a certificate from a CA through a Web browser. After you request a certificate, the CA either sends certificate files by email (public) or gives access to them on a remote host (private). Typically, the CA provides the certificate files listed in [Table 3-6](#).

Table 3-6 SSL Certificate Files

Certificate File	Description
<i>name.crt</i>	The switch certificate.
<i>nameRoot.crt</i>	The root certificate. Typically, this certificate is already installed in the browser, but if not, you must install it.
<i>nameCA.crt</i>	The CA certificate. It is not necessary to install this, but you can if you want the CA name to be displayed in the browser window.

4. On each switch:
 - a. Install the certificate.
 - b. Activate the certificate.
5. If necessary, install the root certificate to the browser on the management workstation.
6. Add the root certificate to the Java Plug-in keystore on the management workstation.

Choosing a Certificate Authority

To ease maintenance and allow secure out-of-band communication between switches, consider using one certificate authority (CA) to sign all management certificates for a fabric. If you use different CAs, management services operate correctly, but the Web Tools Fabric Events button is unable to retrieve events for the entire fabric.

Each CA (for example, Verisign or GeoTrust) has slightly different requirements; for example, some generate certificates based on IP address, while others require an FQDN, and most require a 1024-bit public/private key while some might accept a 2048-bit key. Consider your fabric configuration, check CA Web sites for requirements, and gather all the information that the CA requires.

Generating a Public/Private Key

Perform this procedure on each switch:

1. Connect to the switch and log in as admin.
2. Enter this command to generate a public/private key pair:

```
switch:admin> seccertutil genkey
```

The system reports that this process will disable secure protocols, delete any existing CSR, and delete any existing certificates.

3. Respond to the prompts to continue and select the key size:

```
Continue (yes, y, no, n): [no] y
Select key size [1024 or 2048]: 1024
Generating new rsa public/private key pair
Done.
```

Because CA support for the 2048-bit key size is limited, you should select 1024 in most cases.

Generating and Storing a CSR

After generating a public/private key (see [“Generating a Public/Private Key,”](#) earlier), perform this procedure on each switch:

1. Connect to the switch and log in as admin.
2. Enter this command:

```
switch:admin> seccertutil gencsr
```

3. Enter the requested information:

```
Country Name (2 letter code, eg, US):US
State or Province Name (full name, eg, California):California
Locality Name (eg, city name):San Jose
Organization Name (eg, company name):Brocade
Organizational Unit Name (eg, department name):Eng
Common Name (Fully qualified Domain Name, or IP address): 192.1.2.3
Generating CSR, file name is: 192.1.2.3.csr
Done.
```

Your CA might require specific codes for Country, State or Province, Locality, Organization, and Organizational Unit names. Make sure that your spelling is correct and matches the CA requirements. If the CA requires that the Common Name be specified as an FQDN, make sure that the fully qualified domain name is set on the domain name server.

4. Enter this command to store the CSR:

```
switch:admin> seccertutil export
```

5. Enter the requested information:

```
Select protocol [ftp or scp]: ftp
Enter IP address: 192.1.2.3
Enter remote directory: path_to_remote_directory
Enter Login Name: your account
Enter Password: your password
Success: exported CSR.
```

If you are set up for secure file copy protocol, you can select it; otherwise, select ftp. Enter the IP address of the switch on which you generated the CSR. Enter the remote directory name of the FTP server to which the CSR is to be sent. Enter your account name and password on the server.

Obtaining Certificates

Check the instructions on the CA Web site; then, perform this procedure for each switch:

1. Generate and store the CSR as described in [“Generating and Storing a CSR” on page 3-20](#).
2. Open a Web browser window on the management workstation and go to the CA Web site. Follow the instructions to request a certificate. Locate the area in the request form into which you are to paste the CSR.
3. Through a telnet window, connect to the switch and log in as admin.
4. Enter this command:

```
switch:admin> seccertutil showcsr
```

The contents of the CSR is displayed.

5. Locate the section that begins with “BEGIN CERTIFICATE REQUEST” and ends with “END CERTIFICATE REQUEST”.

6. Copy and paste this section (including the BEGIN and END lines) into the area provided in the request form; then, follow the instructions to complete and send the request.

It might take several days to receive the certificates. If the certificates arrive by email, save them to an FTP server. If the CA provides access to the certificates on an FTP server, make note of the path name and make sure you have a login name and password on the server.

Installing a Switch Certificate

Perform this procedure on each switch:

1. Connect to the switch and log in as admin.
2. Enter this command:

```
switch:admin> seccertutil import
```

3. Select a protocol, enter the IP address of the host on which the switch certificate is saved, and enter your login name and password:

```
Select protocol [ftp or scp]: ftp
Enter IP address: 192.10.11.12
Enter remote directory: path_to_remote_directory
Enter certificate name (must have ".crt" suffix):192.1.2.3.crt
Enter Login Name: your_account
Enter Password: *****
Success: imported certificate [192.1.2.3.crt].
To use this certificate, run the configure command to activate it
```

The certificate downloads to the switch.

Activating a Switch Certificate

Enter the **configure** command and respond to the prompts that apply to SSL certificates:

SSL attributes	Enter yes .
Certificate File	Enter the name of the switch certificate file: for example, 192.1.2.3.crt .
CA Certificate File	If you want the CA name to be displayed in the browser window, enter the name of the CA certificate file; otherwise, skip this prompt.
Select length of crypto key	Enter the encryption key length (40 , 56 , or 128).
HTTP attributes	Enter yes .
Secure HTTP enabled	Enter yes .

Example

```

Configure...
System services (yes, y, no, n): [no]
  ssl attributes (yes, y, no, n): [no] yes
Certificate File. (filename or none): [10.33.13.182.crt] 192.1.2.3.crt
  CA Certificate File. (filename or none): [none]
  Select length of crypto key.
    (Valid values are 40, 56, and 128.): (40..128) [128]
http attributes (yes, y, no, n): [no] yes
HTTP Enabled (yes, y, no, n): [yes] no
  Secure HTTP Enabled (yes, y, no, n): [no] yes

```

After you exit the **configure** command, the HTTP daemon restarts automatically to handle HTTPS requests.

Configuring the Browser

The root certificate might already be installed on your browser, but if not, you must install it. To see whether it is already installed, check the certificate store on your browser.

The next procedures are guides for installing root certificates to Internet Explorer and Mozilla browsers. For more detailed instructions, refer to the documentation that came with the certificate.

To check and install root certificates on Internet Explorer

1. From the browser Tools menu, select **Internet Options**.
2. Click the **Content** tab.
3. Click **Certificates**.
4. Click the various tabs and scroll the lists to see if the root certificate is listed. If it is listed, you do not need to install it, forgo the remainder of this procedure.
5. If the certificate is not listed, click **Import**.
6. Follow the instructions in the Certificate Import wizard to import the certificate.

To check and install root certificates on Mozilla

1. From the browser Edit menu, select **Preferences**.
2. In the left pane of the Preferences window, expand the **Privacy & Security** list and select **Certificates**.
3. In the right pane, click **Manage Certificates**.
4. In the next window, click the **Authorities** tab.
5. Scroll the authorities list to see if the root certificate is listed. (For example, its name might have the form *nameRoot.crt*.) If it is listed, you do not need to install it; forgo the remainder of this procedure.
6. If the certificate is not listed, click **Import**.
7. Browse to the certificate location and select the certificate. (For example, select *nameRoot.crt*.)
8. Click **Open** and follow the instructions to import the certificate.

Installing a Root Certificate to the Java Plug-in

For information on Java requirements, refer to “[Browser and Java Support](#)” on page 3-18.

This procedure is a guide for installing a root certificate to the Java Plug-in on the management workstation. If the root certificate is not already installed to the plug-in, you should install it. For more detailed instructions, refer to the documentation that came with the certificate and to the Sun Microsystems Web site (www.sun.com).

1. Copy the root certificate file from its location on the FTP server to the Java Plug-in bin. For example, the bin location might be:

```
C: \program files\java\j2re1.4.2_03\bin
```

2. Open a Command Prompt window and change directory to the Java Plug-in bin.
3. Enter the **keytool** command and respond to the prompts:

```
C:\Program Files\Java\j2re1.4.2_03\bin> keytool -import -alias RootCert -file
RootCert.crt -keystore ..\lib\security\RootCerts
Enter keystore password: changeit
Owner: CN=Brocade, OU=Software, O=Brocade Communications, L=San Jose,
ST=California, C=US
Issuer: CN=Brocade, OU=Software, O=Brocade Communications, L=San Jose,
ST=California, C=US
Serial number: 0
Valid from: Thu Jan 15 16:27:03 PST 2004 until: Sat Feb 14 16:27:03 PST 2004
Certificate fingerprints:
    MD5: 71:E9:27:44:01:30:48:CC:09:4D:11:80:9D:DE:A5:E3
    SHA1: 06:46:C5:A5:C8:6C:93:9C:FE:6A:C0:EC:66:E9:51:C2:DB:E6:4F:A1
Trust this certificate? [no]: yes
Certificate was added to keystore
```

In the example, **changeit** is the default password and RootCert is an example root certificate name.

Displaying and Deleting Certificates

[Table 3-7](#) summarizes the commands for displaying and deleting certificates. For details on the commands, refer to the *Fabric OS Command Reference Manual*.

Table 3-7 Commands for Displaying and Deleting SSL Certificates

Command	Description
secCertUtil show	Displays the state of the SSL key and a list of installed certificates
secCertUtil show <i>filename</i>	Displays the contents of a specific certificate
secCertUtil showcsr	Displays the contents of a CSR
secCertUtil delete <i>filename</i>	Deletes a specified certificate
secCertUtil delcsr	Deletes a CSR

Troubleshooting Certificates

If you receive messages in the browser or in a pop-up window when logging in to the target switch using HTTPS, refer to [Table 3-8](#).

Table 3-8 SSL Messages and Actions

Message	Action
The page cannot be displayed	The SSL certificate is not installed correctly or HTTPS is not enabled correctly. Make sure that the certificate has not expired, that HTTPS is enabled, and that certificate file names are configured correctly.
The security certificate was issued by a company you have not chosen to trust....	The certificate is not installed in the browser. Install it as described in “Configuring the Browser” on page 3-23 .
The security certificate has expired or is not yet valid	Either the certificate file is corrupted or it needs to be updated. Click View Certificate to verify the certificate content. If it is corrupted or out of date, obtain and install a new certificate.
The name on the security certificate is invalid or does not match the name of the site file	The certificate is not installed correctly in the Java Plug-in. Install it as described in “Installing a Root Certificate to the Java Plug-in” on page 3-24 .
This page contains both secure and nonsecure items. Do you want to display the nonsecure items?	Click No in this pop-up window. The session opens with a closed lock icon on the lower-right corner of the browser, indicating an encrypted connection.

Configuring for SNMP

You can configure for the automatic transmission of Simple Network Management Protocol (SNMP) information to management stations. SNMPv3 and SNMPv1 are supported.

The configuration process involves configuring the SNMP agent and configuring SNMP traps. The following commands are used in the process:

- Use the **configure** command to set the security level. You can specify no security, authentication only, or authentication and privacy.
- Use the **snmpConfig** command to configure the SNMP agent and traps for SNMPv3 or SNMPv1 configurations.
- If necessary for backward compatibility, you can use these legacy commands to configure for SNMP v1:
 - Use the **agtCfgShow**, **agtCfgset**, and **agtCfgDefault** commands to configure the SNMPv1 agent.
 - Use the **snmpMibCapSet** command to filter at the trap level and the **snmpMibCapShow** command to display the trap filter values.

The SNMP trap configuration specifies the MIB trap elements to be used to send information to the SNMP management station. There are two main MIB trap choices:

- Brocade-specific MIB trap

Associated with the Brocade-specific SilkWorm MIB (SW-MIB), this MIB monitors SilkWorm switches specifically.

- FibreAlliance MIB trap

Associated with the FibreAlliance MIB (FA-MIB), this MIB manages SAN switches and devices from any company that complies with FibreAlliance specifications.

If you use both SW-MIB and FA-MIB, you might receive duplicate information. You can disable the FA-MIB, but not the SW-MIB.

You can also use these additional MIBs and their associated traps:

- FICON-MIB (for FICON environments)
- HA-MIB (for SilkWorm 12000 and 24000 models)
- SW-EXTTRAP

Includes the swSsn (Software Serial Number) as a part of Brocade SW traps. It is also used in conjunction with the legacy SilkWorm 6400 integrated fabrics product to provide detailed group information for a particular trap.

For information on Brocade MIBs, refer to the *Fabric OS MIB Reference Manual*.

For information on the specific commands used in these procedures, refer to online help or to the *Fabric OS Command Reference Manual*.

Setting the Security Level

Use the **configure** command to set the security level (called “SNMP attributes”). You can specify no security, authentication only, or authentication and privacy. For example, to configure for authentication and privacy:

```
switch:admin> configure
```

```
Not all options will be available on an enabled switch.
To disable the switch, use the "switchDisable" command.
```

```
Configure...
```

```
System services (yes, y, no, n): [no]
ssl attributes (yes, y, no, n): [no]
http attributes (yes, y, no, n): [no]
snmp attributes (yes, y, no, n): [no] y
```

```
Select SNMP Security Level:
(0 = No security, 1 = Authentication only, 2 = Authentication and Privacy):
(0..2) [0] 2
```

Using the snmpConfig Command

Use the **snmpConfig --set** command to change either the SNMPv3 or SNMPv1 configuration. You can also change access control, MIB capability, and system group.

Sample SNMPv3 configuration

```
switch:admin> snmpconfig --set snmpv3

SNMPv3 user configuration:
User (rw): [snmpadmin1] adminuser
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3] 1
New Auth Passwd:
Verify Auth Passwd:
Priv Protocol [DES(1)/noPriv(2)]: (1..2) [2] 1
New Priv Passwd:
Verify Priv Passwd:
User (rw): [snmpadmin2] shauser
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3] 2
New Auth Passwd:
Verify Auth Passwd:
Priv Protocol [DES(1)/noPriv(2)]: (1..2) [2] 1
New Priv Passwd:
Verify Priv Passwd:
User (rw): [snmpadmin3] nosec
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)]: (2..2) [2]
User (ro): [snmpuser1]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (3..3) [3]
Priv Protocol [DES(1)/noPriv(2)]: (2..2) [2]
User (ro): [snmpuser2]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (3..3) [3]
Priv Protocol [DES(1)/noPriv(2)]: (2..2) [2]
User (ro): [snmpuser3]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (3..3) [3]
Priv Protocol [DES(1)/noPriv(2)]: (2..2) [2]

SNMPv3 trap recipient configuration:
Trap Recipient's IP address in dot notation: [0.0.0.0] 192.168.45.90
UserIndex: (1..6) [1]
Trap recipient Severity level : (0..5) [0] 4
Trap Recipient's IP address in dot notation: [0.0.0.0] 192.168.45.92
UserIndex: (1..6) [2]
Trap recipient Severity level : (0..5) [0] 2
Trap Recipient's IP address in dot notation: [0.0.0.0]
Trap Recipient's IP address in dot notation: [0.0.0.0]
Trap Recipient's IP address in dot notation: [0.0.0.0]
Trap Recipient's IP address in dot notation: [0.0.0.0]
Committing configuration...done.
```

Sample SNMPv1 configuration

```
switch:admin> snmpconfig --set snmpv1

SNMP community and trap recipient configuration:
Community (rw): [Secret C0de] admin
Trap Recipient's IP address in dot notation: [0.0.0.0] 10.32.225.1
Trap recipient Severity level : (0..5) [0] 1
Community (rw): [OrigEquipMfr]
Trap Recipient's IP address in dot notation: [10.32.225.2]
Trap recipient Severity level : (0..5) [1]
Community (rw): [private]
Trap Recipient's IP address in dot notation: [10.32.225.3]
Trap recipient Severity level : (0..5) [2]
Community (ro): [public]
Trap Recipient's IP address in dot notation: [10.32.225.4]
Trap recipient Severity level : (0..5) [3]
Community (ro): [common]
Trap Recipient's IP address in dot notation: [10.32.225.5]
Trap recipient Severity level : (0..5) [4]
Community (ro): [FibreChannel]
Trap Recipient's IP address in dot notation: [10.32.225.6]
Trap recipient Severity level : (0..5) [5]
Committing configuration...done.
```

Sample accessControl configuration

```
switch:admin> snmpconfig --set accessControl

SNMP access list configuration:
Access host subnet area in dot notation: [0.0.0.0] 192.168.0.0
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0] 10.32.148.0
Read/Write? (true, t, false, f): [true] f
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0] 10.33.0.0
Read/Write? (true, t, false, f): [true] f
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Committing configuration...done.
```

Sample mibCapability configuration

```
switch:admin> snmpconfig --show mibCapability
FA-MIB: YES
FICON-MIB: YES
HA-MIB: YES
SW-TRAP: YES
    swFCPortScn: YES
    swEventTrap: YES
    swFabricWatchTrap: YES
    swTrackChangesTrap: NO
FA-TRAP: YES
    connUnitStatusChange: YES
    connUnitEventTrap: NO
    connUnitSensorStatusChange: YES
    connUnitPortStatusChange: YES
SW-EXTTRAP: NO
FICON-TRAP: NO
HA-TRAP: YES
    fruStatusChanged: YES
    cpStatusChanged: YES
    fruHistoryTrap: NO
```

Sample systemGroup configuration (default)

```
switch:admin> snmpconfig --default systemGroup
*****
This command will reset the agent's system group configuration back to factory
default
*****
        sysDescr = Fibre Channel Switch
        sysLocation = End User Premise
        sysContact = Field Support
        authTraps = 0 (OFF)

*****
Are you sure? (yes, y, no, n): [no] y
```

Using Legacy Commands for SNMPv1

You should use the **snmpConfig** command to configure the SNMPv1 agent and traps (refer to [“Using the snmpConfig Command” on page 3-27](#)). However, if necessary for backward compatibility, you can choose to use legacy commands.

Sample SNMP agent configuration information

```
switch:admin> agtcfghshow
Current SNMP Agent Configuration
    Customizable MIB-II system variables:
        sysDescr = FC Switch
        sysLocation = End User Premise
        sysContact = Field Support.
        authTraps = 1 (ON)

SNMPv1 community and trap recipient configuration:
Community 1: Secret C0de (rw)
    Trap recipient: 192.168.1.51
    Trap recipient Severity level: 4
Community 2: OrigEquipMfr (rw)
    Trap recipient: 192.168.1.26
    Trap recipient Severity level: 0
Community 3: private (rw)
    No trap recipient configured yet
Community 4: public (ro)
    No trap recipient configured yet
Community 5: common (ro)
    No trap recipient configured yet
Community 6: FibreChannel (ro)
    No trap recipient configured yet

SNMP access list configuration:
Entry 0: Access host subnet area 192.168.64.0 (rw)]
Entry 1: No access host configured yet
Entry 2: No access host configured yet
Entry 3: No access host configured yet
Entry 4: No access host configured yet
Entry 5: No access host configured yet
```

Sample modification of the SNMP configuration values

```
switch:admin> agtcfgset

Customizing MIB-II system variables ...

At each prompt, do one of the followings:
  o <Return> to accept current value,
  o enter the appropriate new value,
  o <Control-D> to skip the rest of configuration, or
  o <Control-C> to cancel any change.

To correct any input mistake:
<Backspace> erases the previous character,
<Control-U> erases the whole line,
sysDescr: [FC Switch]
sysLocation: [End User Premise]
sysContact: [Field Support.]
authTrapsEnabled (true, t, false, f): [true]

SNMP community and trap recipient configuration:
Community (rw): [Secret C0de]
Trap Recipient's IP address in dot notation: [192.168.1.51]
Trap recipient Severity level : (0..5) [0] 3
Community (rw): [OrigEquipMfr]
Trap Recipient's IP address in dot notation: [192.168.1.26]
Trap recipient Severity level : (0..5) [0]
Community (rw): [private]
Trap Recipient's IP address in dot notation: [0.0.0.0] 192.168.64.88
Trap recipient Severity level : (0..5) [0] 1
Community (ro): [public]
Trap Recipient's IP address in dot notation: [0.0.0.0]
Community (ro): [common]
Trap Recipient's IP address in dot notation: [0.0.0.0]
Community (ro): [FibreChannel]
Trap Recipient's IP address in dot notation: [0.0.0.0]

SNMP access list configuration:
Access host subnet area in dot notation: [0.0.0.0] 192.168.64.0
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Committing configuration...done.
value = 1 = 0x1
```

Sample reset of the SNMP agent configuration to default values

```

switch:admin> agtcfgdefault
*****
This command will reset the agent's configuration back to factory default
*****

Current SNMP Agent Configuration
Customizable MIB-II system variables:
    sysDescr = Fibre Channel Switch.
    sysLocation = End User Premise
    sysContact = sweng
    authTraps = 0 (OFF)
SNMPv1 community and trap recipient configuration:
Community 1: Secret C0de (rw)
    Trap recipient: 192.168.15.41
    Trap recipient Severity level: 4
Community 2: OrigEquipMfr (rw)
    No trap recipient configured yet
Community 3: private (rw)
    No trap recipient configured yet
Community 4: public (ro)
    No trap recipient configured yet
Community 5: common (ro)
    No trap recipient configured yet
Community 6: FibreChannel (ro)
    No trap recipient configured yet
SNMP access list configuration:
Entry 0: Access host subnet area 192.168.64.0 (rw)]
Entry 1: No access host configured yet
Entry 2: No access host configured yet
Entry 3: No access host configured yet
Entry 4: No access host configured yet
Entry 5: No access host configured yet
*****

Are you sure? (yes, y, no, n): [no] y
Committing configuration...done.
agent configuration reset to factory default
Current SNMP Agent Configuration
Customizable MIB-II system variables:
    sysDescr = Fibre Channel Switch.
    sysLocation = End User Premise
    sysContact = Field Support.
    authTraps = 0 (OFF)
SNMPv1 community and trap recipient configuration:
Community 1: Secret C0de (rw)
    No trap recipient configured yet
Community 2: OrigEquipMfr (rw)
    No trap recipient configured yet
Community 3: private (rw)
    No trap recipient configured yet
Community 4: public (ro)
    No trap recipient configured yet
Community 5: common (ro)
    No trap recipient configured yet
Community 6: FibreChannel (ro)
    No trap recipient configured yet
(output truncated)

```

Sample modification of the options for configuring SNMP MIB traps

```
switch:admin> snmpmibcapset
The SNMP Mib/Trap Capability has been set to support
FE-MIB
SW-MIB
FA-MIB
FA-TRAP
FA-MIB (yes, y, no, n): [yes]
FICON-MIB (yes, y, no, n): [no] y
HA-MIB (yes, y, no, n): [no] y
SW-TRAP (yes, y, no, n): [no] y
  swFCPortScn (yes, y, no, n): [no]
  swEventTrap (yes, y, no, n): [no]
  swFabricWatchTrap (yes, y, no, n): [no]
  swTrackChangesTrap (yes, y, no, n): [no]
FA-TRAP (yes, y, no, n): [yes]
  connUnitStatusChange (yes, y, no, n): [no]
  connUnitEventTrap (yes, y, no, n): [no]
  connUnitSensorStatusChange (yes, y, no, n): [no]
  connUnitPortStatusChange (yes, y, no, n): [no]
SW-EXTTRAP (yes, y, no, n): [no] y
FICON-TRAP (yes, y, no, n): [no] y
  linkRNIDDeviceRegistration (yes, y, no, n): [no]
  linkRNIDDeviceDeRegistration (yes, y, no, n): [no]
  linkLIRRLListenerAdded (yes, y, no, n): [no]
  linkLIRRLListenerRemoved (yes, y, no, n): [no]
  linkRLIRFailureIncident (yes, y, no, n): [no]
HA-TRAP (yes, y, no, n): [no] y
  fruStatusChanged (yes, y, no, n): [no]
  cpStatusChanged (yes, y, no, n): [no]
  fruHistoryTrap (yes, y, no, n): [no]
Avoid-Duplicate-TRAP (yes, y, no, n): [no] y
switch:admin>
```

The following **snmpMibCapSet** parameters for FA-TRAP appear in the preceding example:

- **connUnitStatusChange** indicates that the overall status of the connectivity unit has changed. Its variables are:
 - **connUnitStatus**: the status of the connection unit
 - **connUnitState**: the state of the connection unit
- **connUnitEventTrap** indicates that the connectivity unit has generated an event. Its variables are:
 - **connUnitEventId**: the internal event ID
 - **connUnitEventType**: the type of this event
- **connUnitEventObject** is used with the **connUnitEventType** to identify the object to which the event refers.
- **connUnitEventDescr** is the description of the event.
- **connUnitSensorStatusChange** indicates that the status of the sensor associated with the connectivity unit has changed.
- **connUnitSensorStatus** is the status indicated by the sensor.
- **connUnitPortStatusChange** indicates that the status of the sensor associated with the connectivity unit has changed.
- **connUnitPortStatus** shows overall protocol status for the port.
- **connUnitPortState** shows the user-specified state of the port hardware.

Sample view of the SNMP MIB trap setup

```
switch:admin> snmpmibcapshow
FA-MIB: YES
FICON-MIB: YES
HA-MIB: YES
SW-TRAP: YES
    swFCPortScn: YES
    swEventTrap: YES
    swFabricWatchTrap: YES
    swTrackChangesTrap: YES
FA-TRAP: YES
SW-EXTTRAP: YES
HA-TRAP: YES
    fruStatusChanged: YES
    cpStatusChanged: YES
    fruHistoryTrap: YES
```

Configuring Secure File Copy

You can use the **configure** command to specify that secure file copy (scp) be used for configuration uploads and downloads.

Example

```
switch:admin> configure

Not all options will be available on an enabled switch.
To disable the switch, use the "switchDisable" command.

Configure...

System services (yes, y, no, n): [no] n
ssl attributes (yes, y, no, n): [no] n
http attributes (yes, y, no, n): [no] n
snmp attributes (yes, y, no, n): [no] n
rpd attributes (yes, y, no, n): [no] n
cfgload attributes (yes, y, no, n): [no] y

    Enforce secure config Upload/Download (yes, y, no, n): [no] y
switch:admin>
```

Setting the Boot PROM Password

The boot PROM password provides an additional layer of security by protecting the boot PROM from unauthorized use. Setting a recovery string for the boot PROM password enables you to recover a lost boot PROM password by contacting your switch service provider. Without the recovery string, a lost boot PROM password cannot be recovered.

You should set the boot PROM password and the recovery string on all switches, as described in [“With a Recovery String” on page 3-35](#). If your site procedures dictate that you set the boot PROM password without the recovery string, refer to [“Without a Recovery String” on page 3-37](#).



Note

SilkWorm 3016: This model does not have a serial port; therefore, boot PROM procedures do not apply.

With a Recovery String

To set the boot PROM password with a recovery string, refer to the section that applies to your switch model.



Note

Setting the boot PROM password requires accessing the boot prompt, which stops traffic flow through the switch until the switch is rebooted. You should perform this procedure during a planned down time.

SilkWorm 200E, 3250, 3850, 3900, 4012, and 4100

Follow this procedure to set the boot PROM password with a recovery string:

1. Connect to the serial port interface as described in [“To connect through the serial port” on page 2-2](#).
2. Reboot the switch.
3. Press **ESC** within four seconds after the message “Press escape within 4 seconds...” displays.

The following options are available:

Option	Description
1 Start system.	Continues the system boot process.
2 Recovery password.	Lets you set the recovery string and the boot PROM password.
3 Enter command shell.	Provides access to boot parameters.

4. Enter **2**.

If no password was previously set, the following message displays:

```
Recovery password is NOT set. Please set it now.
```

If a password was previously set, the following messages display:

```
Send the following string to Customer Support for password recovery:
afHTpyLsDolPz0Pk5GzhIw==
Enter the supplied recovery password.
Recovery Password:
```

5. Enter the recovery password (string).

The recovery string must be between 8 and 40 alphanumeric characters. A random string that is 15 characters or longer is recommended for higher security. The firmware prompts for this password only once. It is not necessary to remember the recovery string because it is displayed the next time you enter the command shell.

The following prompt displays:

```
New password:
```

6. Enter the boot PROM password; then reenter it when prompted. The password must be eight alphanumeric characters (any additional characters are not recorded). Record this password for future use.

The new password is automatically saved (the **saveEnv** command is not required).

7. Reboot the switch.

SilkWorm 12000, 24000, and 48000

The boot PROM and recovery passwords must be set for each CP blade on SilkWorm 12000, 24000, and 48000 directors as follows:

1. Connect to the serial port interface on the standby CP blade, as described in [“To connect through the serial port” on page 2-2](#).
2. Connect to the active CP blade by serial or telnet and enter the **haDisable** command to prevent failover during the remaining steps.
3. **SilkWorm 12000:** Reboot the standby CP blade by pressing the yellow ejector buttons at top and bottom of the CP blade, and then press both ejector handles back towards the switch to lock the blade back into the slot.

SilkWorm 24000 and 48000: Reboot the standby CP blade by sliding the On/Off switch on the ejector handle of the standby CP blade to Off, and then back to On.

4. Press **ESC** within four seconds after the message “Press escape within 4 seconds...” displays.

The following options are available:

Option	Description
1 Start system.	Continues the system boot process.
2 Recovery password.	Lets you set the recovery string and the boot PROM password.
3 Enter command shell.	Provides access to boot parameters.

5. Enter **2**.

If no password was previously set, the following message displays:

```
Recovery password is NOT set. Please set it now.
```

If a password was previously set, the following messages display:

```
Send the following string to Customer Support for password recovery:
afHTpyLsDo1Pz0Pk5GzhIw==
Enter the supplied recovery password.
Recovery Password:
```

6. Enter the recovery password (string).

The recovery string must be between 8 and 40 alphanumeric characters. A random string that is 15 characters or longer is recommended for higher security. The firmware only prompts for this password once. It is not necessary to remember the recovery string because it is displayed the next time you enter the command shell.

The following prompt displays:

```
New password:
```

7. Enter the boot PROM password; then reenter it when prompted. The password must be eight alphanumeric characters (any additional characters are not recorded). Record this password for future use.

The new password is automatically saved (the **saveEnv** command is not required).
8. Connect to the active CP blade by serial or telnet and enter the **haEnable** command to restore high availability; then fail over the active CP blade by entering the **haFailover** command.

Traffic flow through the active CP blade resumes when the failover is complete.
9. Connect the serial cable to the serial port on the new standby CP blade (previously the active CP blade).
10. Repeat [step 2](#) through [step 7](#) for the new standby CP blade (each CP blade has a separate boot PROM password).
11. Connect to the active CP blade by serial or telnet and enter the **haEnable** command to restore high availability.

Without a Recovery String

Although you can set the boot PROM password without also setting the recovery string, it is strongly recommended that you set both the password and the string as described in [“With a Recovery String” on page 3-35](#). If your site procedures dictate that you must set the boot PROM password without the string, follow the procedure that applies to your switch model.



Note

Setting the boot PROM password requires accessing the boot prompt, which stops traffic flow through the switch until the switch is rebooted. You should perform this procedure during a planned down time.

SilkWorm 200E, 3250, 3850, 3900, 4012, and 4100

Follow this procedure to set the boot PROM password without a recovery string:

1. Create a serial connection to the switch as described in [“To connect through the serial port” on page 2-2](#).
2. Reboot the switch by entering the **reboot** command.
3. Press **ESC** within four seconds after the message “Press escape within 4 seconds...” displays.

The following options are available:

Option	Description
1 Start system.	Continues the system boot process.
2 Recovery password.	Lets you set the recovery string and the boot PROM password.
3 Enter command shell.	Provides access to boot parameters.

4. Enter **3**.
5. At the shell prompt, enter the **passwd** command.



Note

The **passwd** command only applies to the boot PROM password when it is entered from the boot interface.

6. Enter the boot PROM password at the prompt, then reenter it when prompted. The password must be eight alphanumeric characters (any additional characters are not recorded). Record this password for future use.
7. Enter the **saveEnv** command to save the new password.
8. Reboot the switch by entering the **reset** command.

SilkWorm 12000, 24000, and 48000

On the SilkWorm 12000, 24000, and 48000, set the password on the standby CP blade, fail over, and then set the password on the previously active (now standby) CP blade to minimize disruption to the fabric:

1. Determine the active CP blade by opening a telnet session to either CP blade, connecting as admin, and entering the **haShow** command.
2. Connect to the active CP blade by serial or telnet and enter the **haDisable** command to prevent failover during the remaining steps.
3. Create a serial connection to the standby CP blade as described in [“To connect through the serial port” on page 2-2](#).
4. **SilkWorm 12000:** Reboot the standby CP blade by pressing the yellow ejector buttons at top and bottom of the CP blade, then pressing both ejector handles back towards the switch to lock the blade back into the slot.

SilkWorm 24000 and 48000: Reboot the standby CP blade by sliding the On/Off switch on the ejector handle of the standby CP blade to Off, and then back to On.

This causes the blade to reset.

5. Press **ESC** within four seconds after the message “Press escape within 4 seconds...” displays.

The following options are available:

Option	Description
1 Start system.	Continues the system boot process.
2 Recovery password.	Lets you set the recovery string and the boot PROM password.
3 Enter command shell.	Provides access to boot parameters.

6. Enter **3**.
7. Enter the **passwd** command at the shell prompt.



Note

The **passwd** command applies only to the boot PROM password when it is entered from the boot interface.

8. Enter the boot PROM password at the prompt; then reenter it when prompted. The password must be eight alphanumeric characters (any additional characters are not recorded). Record this password for future use.
9. Enter the **saveEnv** command to save the new password.
10. Reboot the standby CP blade by entering the **reset** command.
11. Connect to the active CP blade by serial or telnet and enter the **haEnable** command to restore high availability; then fail over the active CP blade by entering the **haFailover** command.
Traffic resumes flowing through the newly active CP blade after it has completed rebooting.
12. Connect the serial cable to the serial port on the new standby CP blade (previously the active CP blade).
13. Repeat [step 3](#) through [step 10](#) for the new standby CP blade.
14. Connect to the active CP blade by serial or telnet and enter the **haEnable** command to restore high availability.

Recovering Forgotten Passwords

Passwords can be recovered as follows:

- If you know the root password, you can use this procedure to recover the user, admin, and factory passwords:
 1. Open a CLI connection (serial or telnet) to the switch. If secure mode is enabled, connect to the primary FCS switch.
 2. Log in as root.
 3. Enter the command for the type of password that was lost:


```
passwd user
passwd admin
passwd factory
```
 4. Enter the requested information at the prompts.

- To recover a lost root password, contact your switch service provider.
- To recover a lost boot PROM password, contact your switch service provider. You must have previously set a recovery string to recover the boot PROM password.

Maintaining Configurations and Firmware

This chapter contains procedures for maintaining switch configurations and for installing firmware.

Maintaining Configurations

It is important to maintain consistent configuration settings on all switches in the same fabric, because inconsistent parameters (such as inconsistent PID formats) can cause fabric segmentation. As part of standard configuration maintenance procedures, it is recommended that you back up all important configuration data for every switch on a host computer server for emergency reference.

The following sections contain procedures for basic switch configuration maintenance.

Displaying Configuration Settings

The switch configuration file comprises three sections, and is organized as follows:

- The Boot Parameters section contains variables such as the switch's name and IP address.
- The Licenses section lists the licenses that are active on the switch.
- The Chassis Configuration section contains configuration variables such as diagnostic settings, fabric configuration settings, and SNMP settings.
- The Configuration section contains licensed option configuration parameters.

To display configuration settings, connect to the switch, log in as admin, and enter the **configShow** command at the command line. The configuration settings vary depending on switch model and configuration.

Backing Up a Configuration

In case the configuration is lost or unintentional changes are made, keep a backup copy of the configuration file (if using a SilkWorm 12000 or 24000 with two domains, configuration *files*—one for each logical switch), which is what gets saved when you enter the **configUpload** command. You should keep individual backup files for all switches in the fabric. You should avoid copying configurations from one switch to another.

The following information is *not* saved in a backup:

- dnsconfig information
- passwords

You must have a valid account on the FTP server where the backup file is to be stored.

You can specify the use of secure file copy (scp) during the procedure. For instructions on configuring the use of scp by default, refer to [“Configuring Secure File Copy” on page 3-34](#).

Before beginning, verify that you can reach the FTP server from the switch. Using a telnet connection, save a backup copy of the configuration file from a logical switch to a host computer as follows:

1. Verify that the FTP service is running on the host computer.
2. Connect to the switch and log in as admin.
3. Enter the **configUpload** command. The command becomes interactive and you are prompted for the required information.
4. Respond to the prompts as follows:

<i>Protocol (scp or ftp)</i>	If your site requires the use of Secure Copy, specify scp. Otherwise, specify ftp.
<i>Server Name or IP Address</i>	Enter the name or IP address of the server where the file is to be stored; for example, 192.1.2.3. You can enter a server name if DNS is enabled. For details about the dnsConfig command, refer to the <i>Fabric OS Command Reference Manual</i> .
<i>User name</i>	Enter the user name of your account on the server; for example, “JohnDoe”.
<i>File name</i>	Specify a file name for the backup file; for example, “config.txt”. Absolute path names can be specified using forward slash (/). Relative path names create the file in the user’s home directory on UNIX servers, and in the directory where the FTP server is running on Windows servers.
<i>Password</i>	Enter your account password for the server.

Example

```
switch:admin> configupload
Protocol (scp or ftp) [ftp]: ftp
Server Name or IP Address [host]: 192.1.2.3
User Name [user]: JohnDoe
File Name [config.txt]: /pub/configurations/config.txt
Password: xxxxxx
Upload complete
switch:admin>
```

Restoring a Configuration

Restoring a configuration involves overwriting the configuration on the switch by downloading a previously saved backup configuration file. Perform this procedure during a planned down time.

Make sure that the configuration file you are downloading is compatible with your switch model, because configuration files from other model switches might cause your switch to fail.

You must have a user ID on the FTP server where the backup file is stored.

Use the following procedure:

1. Verify that the FTP service is running on the server where the backup configuration file is located.
2. Connect to the switch and log in as admin.
3. Disable the switch by entering the **switchDisable** command.
4. Enter the **configDownload** command.

The command becomes interactive and you are prompted for the required information.

5. Respond to the prompts as follows:

<i>Protocol (scp or ftp)</i>	If your site requires the use of Secure Copy, specify scp. Otherwise, specify ftp.
<i>Server Name or IP Address</i>	Enter the name or IP address of the server where the file is stored; for example, 192.1.2.3. You can enter a server name if DNS is enabled.
<i>User name</i>	Enter the user name of your account on the server; for example, "JohnDoe".
<i>File name</i>	Specify the full path name of the backup file; for example, "/pub/configurations/config.txt".
<i>Password</i>	Enter your account password for the server.

6. At the "Do you want to continue [y/n]" prompt, enter **y**.
7. Wait for the configuration to be restored.
8. When the process is finished, enter the **switchEnable** command.

```
switch:admin> configdownload
Protocol (scp or ftp) [ftp]: ftp
Server Name or IP Address [host]: 192.1.2.3
User Name [user]: JohnDoe
File Name [config.txt]: /pub/configurations/config.txt
Password: xxxxxx

*** CAUTION ***

This command is used to download a backed-up configuration
for a specific switch. If using a file from a different
switch, this file's configuration settings will override
any current switch settings. Downloading a configuration
file, which was uploaded from a different type of switch,
may cause this switch to fail.

Do you want to continue [y/n]: y
download complete..
switch:admin> switchenable
```



Note

Because some configuration parameters require a reboot to take effect, after you download a configuration file you must reboot to be sure that the parameters are enabled. Before the reboot, this type of parameter is listed in the configuration file, but it is not effective until after the reboot.

Restoring Configurations in a FICON Environment

If the switch is operating in a FICON CUP environment, and the ASM (active=saved) bit is set on, then the switch ignores the IPL file downloaded when you restore a configuration. [Table 4-1](#) describes this behavior in more detail.

Table 4-1 Backup and Restore in a FICON CUP Environment

ASM bit	Command	Description
on or off	configupload	All the files saved in file access facility are uploaded to the management workstation. A section in the uploaded configuration file labeled FICON_CUP is in an encoded format.
on	configdownload	Files saved on the switch that are also present in the FICON_CUP section of the configuration file are overwritten. Files in the FICON section of configuration file that are not currently present on the switch are saved. The IPL file is not replaced, because active=saved mode is on. A warning message is displayed in the syslog to warn that the IPL file is not being overwritten.
off	configdownload	Files saved on the switch that are also present in the FICON_CUP section of the configuration file are overwritten. Files in the FICON section of configuration file that are not currently present on the switch are saved. The IPL file is replaced, because active=saved mode is off.

If **fmsmode** is enabled in a configuration file, but is disabled on the switch, the **configdownload** command fails and displays an error message. This prevents undesirable conditions that could result from enabling **fmsmode** on a switch that does not require it.

Downloading Configurations Across a Fabric

To save time when configuring fabric parameters and software features, you can save a configuration file from one switch and download it to other switches of the same model type, as shown in the following procedure. Avoid downloading configuration files to different model switches, because that can cause the switches to fail.

1. Configure one switch first.
2. Use the **configUpload** command to save the configuration information. Refer to “[Backing Up a Configuration](#)” on page 4-1.
3. Use the **configDownload** command to download it onto each of the remaining switches. Refer to “[Restoring a Configuration](#)” on page 4-2.

Printing Hard Copies of Switch Information

It is recommended that you print a hard copy of all key configuration data, including license key information for every switch, and store it in a safe and secure place for emergency reference. Print out the information from the following commands, and store the printouts in a secure location:

- **configShow**
Displays configuration parameters and setup information, including license information.
- **ipAddrShow**
Displays the IP address.
- **licenseShow**
Displays the license keys you have installed and provides better detail than the license information from the **configShow** command.

Depending on the security procedures of your company, you might also want to keep a record of the user levels and passwords (including any boot ROM passwords) for all switches in the fabric. Access to this sensitive information should be limited.

Maintaining Firmware

This section explains how to obtain and install firmware. Fabric OS v5.0.1 provides nondisruptive firmware installation.

In most cases, you will be *upgrading* firmware; that is, installing a newer firmware version than the one you are currently running. However, some circumstances might require installing an older version; that is, *downgrading* the firmware. The procedures in this section assume that you are upgrading firmware, but they work for downgrading as well, provided the old and new firmware versions are compatible. Also, always reference the latest release notes for updates that may exist regarding downgrades under particular circumstances.

Using the CLI (or Brocade Web Tools), you can upgrade the firmware on one switch at a time. You can use the optionally licensed Brocade Fabric Manager software tool to upgrade firmware simultaneously on multiple switches. For more details on Fabric Manager and other licensed software tools, go to the Brocade Web site at <http://www.brocade.com>.

Obtaining and Unzipping Firmware

Firmware upgrades are available for customers with support service contracts and partners on the Brocade Web site at <http://www.brocade.com>.

At the Brocade Web site, click *Brocade Connect* and follow the instructions to register and download firmware. Partners with authorized accounts can use the *Brocade Partner Network*.

You must unzip the firmware (using the UNIX **tar** or **gzip** command or a Windows unzip program) before you can use the **firmwareDownload** command to update the firmware on your equipment.

When you unpack the downloaded firmware it expands into a directory that is named according to the version of Fabric OS it contains. For example, if you download and unpack Fabric OS v5.0.1.zip, it expands into a directory called v5.0.1. When you use the **firmwaredownload** command, you specify the path to the v5.0.1 directory and append the keyword **release.plist** to the path.

Checking Connected Switches

If the switch to be upgraded is running v4.1.0 firmware (or later), it is recommended that all switches directly connected to it be running versions no earlier than v2.6.1, v3.1.0, or v4.1.0. If some connected switches are running older firmware, upgrade them to *at least* the earliest recommended version (shown in [Table 4-2](#)) before upgrading firmware on your switch.



Note

Please go to <http://www.brocade.com> to view end-of-life policies for Brocade products. Navigate to the **Services and Support** tab, then select **End of Life Information**. End-of-life products are not supported.

Table 4-2 Recommended Firmware

SilkWorm Model ^a	Earliest Recommended Fabric OS Version
200E	v5.0.1
2000 series	v2.6.1
3200, 3600, 3800	v3.1.0
3016, 3250, 3850	v4.2.0
3900	v4.1.0
4012	v5.0.0
4100	v4.4.0
12000	v4.1.0
24000	v4.2.0
48000	v5.0.1

a. During code activation on SilkWorm 3016, 3250, 3850, or 3900 running Fabric OS v4.1.0 or later, data continues to flow between hosts and storage devices; however, fabric services are unavailable for a period of approximately 50-55 seconds. Possible disruption of the fabric can be minimized by ensuring that switches logically adjacent to these models (directly connected via an ISL) are running at the minimum Fabric OS v2.6.1 or later, v3.1.0 or later, or v4.1.0 or later.

If SilkWorm 3016, 3250, 3850, 3900, 4012, or 4100 models are adjacent and you start firmware downloads on them at same time, there might be I/O disruption.

To determine whether you need to upgrade connected switches before upgrading your switch, use the following procedure on each connected switch to display firmware information and build dates.

1. Connect to the switch and log in as admin.
2. Enter the **version** command.

The following information is displayed:

Kernel:	Displays the version of switch kernel operating system.
Fabric OS:	Displays the version of switch Fabric OS.
Made on:	Displays the build date of firmware running in switch.
Flash:	Displays the install date of firmware stored in nonvolatile memory.
BootProm:	Displays the version of the firmware stored in the boot PROM.

About the Download Process

The **firmwareDownload** command downloads unzipped switch firmware from an FTP server to the switch's nonvolatile storage area.

In the SilkWorm 12000, 24000, and 48000 directors, this command (when not using any options flags) by default downloads the firmware image to a standby CP, if there is one, to prevent disruption to application services. This operation depends on high-availability (HA) support. If HA is not available, experienced technicians can upgrade the CPs one at a time, using the **-s** option.



For each non-director-class switch in your fabric, complete all firmware download changes before issuing the **firmwareDownload** command on the next switch to ensure a nondisruptive download.

SilkWorm fixed-port models and each CP blade of the SilkWorm 12000, 24000, and 48000 models have two partitions of nonvolatile storage areas (a primary and a secondary) to store two firmware images. The **firmwareDownload** command always loads the new image into the secondary partition and swaps the secondary partition to be the primary. It then reboots the partition and activates the new image. Finally, it performs the **firmwareCommit** procedure automatically, to copy the new image to the other partition.

Effects of Firmware Changes on Accounts and Passwords

The following table describes what happens to accounts and passwords when you replace the switch firmware with a different version. *Upgrading* means installing a newer version of firmware. *Downgrading* means installing an older version of firmware.

Table 4-3 Effects of Firmware Changes on Accounts and Passwords

Change	First Time	Subsequent Times (After upgrade, then downgrade, then upgrade)
Upgrading	Default accounts and their passwords are preserved.	User-defined and default accounts and their passwords are preserved.
Downgrading	User-defined accounts are no longer valid. Default accounts and their passwords are preserved. If a default account was disabled, it is reenabled after the downgrade.	User-defined and default accounts and their passwords are preserved, including accounts added after the first upgrade.
Upgrading to v3.2.0	(You might upgrade a switch in the fabric as part of “Checking Connected Switches” on page 4-6.) Earlier versions allowed you to change the default account names. You cannot add user-defined accounts until you change the names back to default with the passwdDefault command.	

For more details on older releases of Fabric OS, refer to [“Understanding Legacy Password Behavior” on page D-1.](#)

Considerations for Downgrading Firmware

The following items must be considered before attempting to downgrade to an earlier version of Fabric OS:

- If your fabric is set to the extended edge PID format and you want to downgrade to an older Fabric OS version that does not support extended edge, you must change the PID to a supported format. For more information, refer to [“Configuring the PID Format” on page A-1.](#)
- Downgrading a SilkWorm 24000 that is configured for two domains from Fabric OS v4.4.0 to Fabric OS v4.2.0 is not supported.
- If you are running v4.0.2 firmware on a SilkWorm 3900, you cannot downgrade to earlier versions.

Considerations for FICON CUP Environments

To prevent channel errors during nondisruptive firmware installation, the switch CUP port must be taken offline from all host systems.

Upgrading SilkWorm Switches

SilkWorm 200E, 3016, 3250, 3850, 3900, 4012, and 4100 switches maintain primary and secondary partitions for firmware. The **firmwareDownload** command defaults to an autocommit option that automatically copies the firmware from one partition to the other.

You should not override autocommit under normal circumstances; use the default. Refer to [Appendix F, “Upgrading Firmware in Single CP Mode”](#) for details about overriding the autocommit option.

Optionally, before starting a firmware download, you can connect the switch with a serial console cable to a computer that is running a session capture. The information collected might be useful for troubleshooting.

Summary of the Upgrade Process

The following summary describes the default behavior of the **firmwareDownload** command (without options) on SilkWorm 200E, 3016, 3250, 3850, 3900, 4012, and 4100 models:

1. You enter the **firmwareDownload** command.
2. Fabric OS downloads firmware to the secondary partition.
3. The system performs a high-availability reboot (**haReboot**). After the **haReboot**, the former secondary partition is the primary partition.
4. The system replicates the firmware from the primary to the secondary partition.

Enter the **firmwareDownloadStatus** command to view the firmware process.

SilkWorm 200E, 3016, 3250, 3850, 3900, 4012, and 4100 Upgrade Procedure

The upgrade process first downloads and then commits the firmware to the switch. While the upgrade is proceeding, you can start another telnet session on the switch and observe the upgrade progress if you wish.



Note

After you start the process, do not enter any disruptive commands (such as reboot) that will interrupt the process. The entire firmware download and commit process takes approximately 17 minutes. If there is a problem, wait for the timeout (30 minutes for network problems; 10 minutes for incorrect IP address). Disrupting the process can render the switch inoperable and require you to seek help from Customer Support.

Do not disconnect the switch from power during the process, because the switch could become inoperable upon reboot.

Use this procedure to upgrade firmware for SilkWorm 200E, 3016, 3250, 3850, 3900, 4012, and 4100 switches:

1. Verify that the FTP service is running on the host server and that you have a user ID on that server.
2. Obtain the firmware file from the Brocade Web site at <http://www.brocade.com> and store the file on the FTP server. Verify that the FTP service is running and unpack the compressed files preserving directory structures.
3. Use the **firmwareShow** command to check the current firmware version on connected switches. Upgrade their firmware if necessary before proceeding with upgrading this switch.
Refer to “[Checking Connected Switches](#)” on page 4-6.
4. Connect to the switch and log in as admin.

5. Use the **firmwareShow** command to check the current firmware version of the switch to verify compatibility with the version of firmware you are going to download.



Note

SilkWorm 3016: If you are running Fabric OS v4.2.1 firmware, you cannot downgrade to earlier versions.

SilkWorm 3250 and 3850: If you are running Fabric OS v4.2.0 firmware, you cannot downgrade to earlier versions.

SilkWorm 3900: If you are running Fabric OS v4.0.2 firmware, you cannot downgrade to earlier versions.

6. Enter the **firmwareDownload** command.
7. At the “Do you want to continue [y/n]” prompt, enter **y**.
8. Respond to the prompts as follows:

Server Name or IP Address Enter the name or IP address of the server where the firmware file is stored; for example, 192.1.2.3. You can enter a server name if DNS is enabled.

User name Enter the user name of your account on the server; for example, “JohnDoe”.

File name Specify the full path name of the firmware directory, appended by release.plist; for example, /pub/v5.0.1/release.plist.

Note: For v4.x and v5.x switches only, do not attempt to locate the release.plist file in the top level directory; there is a release.plist file for each platform, and the correct one is automatically selected.

Password Enter your account password for the server.

After the firmware is downloaded, the switch reboots and starts the firmware commit.

9. After the reboot, connect to the switch and log in again as admin.
10. If you want to watch the upgrade progress, enter the **firmwareDownloadStatus** command.
11. After the firmware commit finishes, enter the **firmwareShow** command to display the firmware level for both partitions.

Example

```

switch:admin> firmwaredownload
You can run firmwareDownloadStatus to get the status of this command.
This command will cause the switch to reset and will require that existing
telnet, secure telnet or SSH sessions be restarted.
Do you want to continue [Y]: y
Server Name or IP Address: 192.1.2.3
User Name: JohnDoe
File Name: /pub/v5.0.1/release.plist
Password: xxxxxx
Firmwaredownload has started.

0x8fd (Fabric OS): Switch: 0, Warning SULIB-FWDL_START, 3, Firmwaredownload
command has started.
.
.
.
```

Log in again to view the upgrade progress:

```

switch:admin> firmwaredownloadstatus
[0]: Tue Apr 20 10:32:34 2004
cp0: Firmwaredownload has started.
[1]: Tue Apr 20 10:36:07 2004
cp0: Firmwaredownload has completed successfully.
[2]: Tue Apr 20 10:57:09 2004
cp0: Firmwarecommit has started.
[3]: Tue Apr 20 10:36:07 2004
cp0: Firmwarecommit has completed successfully.
[4]: Tue Apr 20 11:03:28 2004
cp0: Firmwaredownload command has completed successfully.
switch:admin> firmwareshow
Primary partition: v5.0.1
Secondary Partition: v5.0.1
switch:admin>
```

Upgrading SilkWorm Directors

You can download firmware to SilkWorm 12000, 24000, and 48000 directors without disrupting the overall fabric if the two CP blades are installed and fully synchronized. Use the **haShow** command to confirm synchronization. If only one CP blade is powered on, the switch must reboot to activate firmware, which is disruptive to the overall fabric.

If there is an error during the firmware download, the system ensures that the two partitions of a CP blade contain the same version of firmware. However, the two CP blades might contain different versions of firmware; in that event, repeat the firmware download process.

During the upgrade process, the director fails over to its standby CP blade and the IP addresses for the two logical switches move to that CP blade's Ethernet port. This might cause informational ARP address reassignment messages to appear on other switches in the fabric. This is normal behavior, because the association between the IP addresses and MAC addresses has changed.



Caution

To successfully download firmware to a director you must have an active Ethernet connection on *both* CPs.

Summary of the Upgrade Process

The following summary describes the default behavior of the **firmwareDownload** command (without options) on SilkWorm 12000, 24000, and 48000 directors:

1. You enter the **firmwareDownload** command on the active CP blade.
2. The standby CP blade downloads firmware.
3. The standby CP blade reboots and comes up with the new Fabric OS.
4. The active CP blade synchronizes its state with the standby CP blade.
5. The active CP blade forces a failover and reboots to become the standby CP blade.
6. The *new* standby CP blade (the active CP blade before the failover) downloads firmware.
7. The *new* standby CP blade reboots and comes up with the new Fabric OS.
8. The new active CP blade synchronizes its state with the new standby CP blade.
9. The **firmwareCommit** command runs automatically on both CP blades.



Note

After you start the process, do not enter any disruptive commands (such as reboot) that will interrupt the process. The entire firmware download and commit process takes approximately 15 minutes. If there is a problem, wait for the timeout (30 minutes for network problems; 10 minutes for incorrect IP address). Disrupting the process can render the switch inoperable and require you to seek help from Customer Support.

Do not disconnect the switch from power during the process, because the switch could become inoperable upon reboot.

SilkWorm 12000, 24000, and 48000 Upgrade Procedure

SilkWorm 12000 directors have four IP addresses: one for each of the two logical switches (switch 0 and switch 1) and one for each of the two CP blades (CP0 in slot 5 and CP1 in slot 6). The SilkWorm 24000 director in its default configuration has three IP addresses but can be configured for four. The SilkWorm 48000 does not support two domains, hence, you can use only three IP addresses.



Note

By default, the **firmwareDownload** command automatically upgrades both the active CP blade and the standby CP blade. When upgrading a SilkWorm 12000 that is running v4.0.0c or earlier, you must upgrade each CP blade separately, as described in [“To upgrade a single SilkWorm 12000/24000 CP blade” on page F-2](#). (You should not use this procedure under normal circumstances.)

Follow this procedure to upgrade the firmware on SilkWorm 12000, 24000, and 48000 directors:

1. Verify that the FTP service is running on the host server and that you have a user ID on that server.
2. Obtain the firmware file from the Brocade Web site at <http://www.brocade.com> and store the file on the FTP server. Verify that the FTP service is running and unpack the compressed files preserving directory structures.
3. Use the **firmwareShow** command to check the current firmware version on connected switches. Upgrade the firmware, if necessary, before proceeding with upgrading this switch.

Refer to “Checking Connected Switches” on page 4-6.

4. Using a telnet session, connect to the switch and log in as admin.
5. **SilkWorm 12000:** Use the **firmwareShow** command to check the current firmware version of the switch.
6. Enter the **haShow** command to confirm that the two CP blades are synchronized.

CP blades must be synchronized and running Fabric OS v4.1.0 or later to provide a nondisruptive download. If the two CP blades are not synchronized, and the current firmware version is 4.1.0 or later, enter the **haSyncStart** command to synchronize the two CP blades. In the following example, the active CP blade is CP1 and the standby CP blade is CP0:

```
switch:admin> hashow
Local CP (Slot 6, CP1): Active
Remote CP (Slot 5, CP0): Standby
HA Enabled, Heartbeat up, HA State is in Sync
switch:admin>
```

7. Log in to either of the logical switches (sw0 for SilkWorm 48000 and SilkWorm 24000 configured with a single domain).
8. Enter the **firmwareDownload** command.
9. At the “Do you want to continue [y/n]” prompt, enter **y**
10. Respond to the prompts as follows:

<i>Server Name or IP Address</i>	Enter the name or IP address of the server where the firmware file is stored: for example, 192.1.2.3. You can enter a server name if DNS is enabled.
<i>User name</i>	Enter the user name of your account on the server: for example, “JohnDoe”.
<i>File name</i>	Specify the full path name of the firmware directory, appended by release.plist: for example, /pub/v5.0.1/release.plist.
<i>Password</i>	Enter your account password for the server.

The firmware is downloaded to one CP blade at a time, beginning with the standby CP blade. During the process, the active CP blade is failed over. After the firmware is downloaded, a firmware commit starts on both CP blades. The entire firmware download and commit process takes approximately 35 minutes.

11. Optionally, after the failover, connect to the switch and log in again as admin.
12. Using a separate telnet session, enter the **firmwareDownloadStatus** command to monitor the firmware download status.
13. Enter the **firmwareShow** command to display the new firmware versions.

Example

```

switch:admin> firmwaredownload
This command will upgrade both CPs in the switch. If you
what to upgrade a single CP only, please use -s option.

You can run firmwareDownloadStatus to get the status
of this command.

This command will cause the active CP to reset and will
require that existing telnet, secure telnet, or SSH sessions
be restarted.

Do you want to continue [Y]: y
Server Name or IP Address: 192.1.2.3
User Name: JohnDoe
File Name: /pub/v5.0.1/release.plist
Password:*****
FirmwareDownload has started on Standby CP. It may take up to 30 minutes.
Firmwaredownload has completed successfully on Standby CP.
.
.
.
Standby CP reboots.
Standby CP booted up.
Standby CP booted up with new firmware.
cpl: Firmwarecommit has started on both Active and Standby CPs.
cpl: Firmwarecommit has completed successfully on Active CP.
cpl: Firmwaredownload command has completed successfully.
switch:admin>

```

Start a new session to view the upgrade progress:

```

switch:admin> firmwaredownloadstatus
[0]: Tue Apr 20 15:18:56 2003
cp0: Firmwaredownload has started on Standby CP. It may take up to 10 minutes.
[1]: Tue Apr 20 15:24:17 2003
cp0: Firmwaredownload has completed successfully on Standby CP.
[2]: Tue Apr 20 15:24:19 2003
cp0: Standby CP reboots.
[3]: Tue Apr 20 15:27:06 2003
cp0: Standby CP booted up.
[4]: Tue Apr 20 15:29:01 2003
cp1: Active CP forced failover succeeded. Now this CP becomes Active.
[5]: Tue Apr 20 15:29:05 2003
cp1: Firmwaredownload has started on Standby CP. It may take up to 30 minutes.
[6]: Tue Apr 20 15:34:16 2003
cp1: Firmwaredownload has completed successfully on Standby CP.
[7]: Tue Apr 20 15:34:19 2003
cp1: Standby CP reboots.
[8]: Tue Apr 20 15:36:59 2003
cp1: Standby CP booted up with new firmware.
[9]: Tue Apr 20 15:37:04 2003
cp1: Firmwarecommit has started on both Active and Standby CPs.
[10]: Tue Apr 20 15:42:48 2003
cp1: Firmwarecommit has completed successfully on Active CP.
[11]: Tue Apr 20 15:42:49 2003
cp1: Firmwaredownload command has completed successfully.

```

Troubleshooting Firmware Downloads

A firmware download can fail for many reasons, such as a power failure, a failed network connection, a failed FTP server, or an incorrect path to unpacked firmware files. In most cases, the firmware will not be affected. You can make necessary corrections (for example, check the Ethernet cables and check the file path names) and then run the **firmwareDownload** command again.



Note

Under firmware versions earlier than v4.1.0, do not perform a firmware download while the switch is running POST. If a firmware download is attempted on a SilkWorm 12000 director while POST is running, it might fail because the CP blades cannot synchronize with each other.

Enter the **firmwareShow** command to see if both CP blades have the same firmware. In this example, the active CP blade and the standby CP blade have the same version:

```
switch: admin> firmwareshow
Local CP (Slot 6, CP1): Standby
    Primary partition:    v5.0.1d
    Secondary Partition:  v5.0.1d
Remote CP (Slot 5, CP0): Active
    Primary partition:    v5.0.1d
    Secondary Partition:  v5.0.1d

Note: If Local CP and Remote CP have different versions
of firmware, please retry firmwaredownload command.
switch: admin>
```

Decide which firmware version you want to be applied to each CP blade. If you want the version on the standby CP, then enter the **haFailover** command on the active CP. If you want the version from the active CP, then enter the **firmwareDownload -s** command on the standby CP. After entering the **haFailover** command, you must enter the **firmwareDownload -s** command on the new standby CP.

Configuring SilkWorm 12000, 24000, and 48000 Directors

This chapter contains procedures that are specific to SilkWorm 12000, 24000, and 48000 Directors.

Because directors contain interchangeable 16-port blades (32-port blades in the SilkWorm 48000), their procedures differ from those for SilkWorm 200E, 3016, 3250, 3850, 3900, 4012, and 4100 fixed-port switches. For example, fixed-port models identify ports by *domain, port* number, while director models identify ports by *slot/port* number.

Also, because the SilkWorm 12000 director comprises two logical switches (*domains*), and the SilkWorm 24000 and 48000 directors in their default configurations have only one domain (SilkWorm 48000 supports one domain only), procedures for the directors sometimes differ from one another.

For detailed information about the SilkWorm 12000, 24000, and 48000 models, refer to the *SilkWorm 12000 Hardware Reference*, the *SilkWorm 24000 Hardware Reference*, and the *SilkWorm 48000 Hardware Reference*.

For detailed information about migrating a SilkWorm director and supported migration paths, refer to the *SilkWorm Director Blade Support Notes*.

Identifying Ports

SilkWorm 12000, 24000, and 48000 models have slots and can have a variable number of ports within a given domain. Ports are identified by their combined slot number and port number.

There are 10 slots that contain port blades:

- Slot numbers 5 and 6 contain *control processor* blades (CPs).
- Slot numbers 1 through 4 and 7 through 10 contain *port* blades.

On each port blade, there are 16 or 32 ports (counted from the bottom, 0 to 15, or 0 to 31). A particular port must be represented by both slot number (1 through 4 and 7 through 10) and port number (0 through 15).

When you have port blades with different port counts in the same director (for example, 16-port blade and 32-port blades), the area IDs no longer match the port numbers. Following are the port numbering schemes for the SilkWorm 48000:

- For the FC4-16 port blade, ports are numbered from 0 through 15 from bottom to top.
- For the FC-32 port blade, ports are numbered from 0 through 15 from bottom to top on the left set of ports and 16 through 31 from bottom to top on the right set of ports.

The SilkWorm 12000 is divided into two logical switches, where slot 1 through 4 is logical switch 0 (sw0) and slot 7 through 10 is logical switch 1 (sw1). You must be connected to the logical switch that represents the slot where you want to execute a command.

In the SilkWorm 24000 and 48000 default configuration, all the ports are part of a single logical switch. With Fabric OS v4.4.0 and later, you can configure the SilkWorm 24000 as two logical switches (*domains*).

The following sections tell how to identify ports on SilkWorm 12000, 24000, and 48000 models, and how to identify ports for zoning commands.

By Slot and Port Number

The port number is a number assigned to an external port to give it a unique identifier in a switch.

To select a specific port in the SilkWorm 12000, 24000, and 48000 models, you must identify both the slot number and the port number using the format *slot number/port number*. No spaces are allowed between the slot number, the slash (/), and the port number.

The following example shows how to enable port 4 on a blade in slot 2:

```
switch:admin> portenable 2/4
```

By Port Area ID

Zoning commands require that you specify ports using the area ID method. In Fabric OS v4.0.0 and later, each port on a particular domain is given a unique area ID. The relationship between the port number and area ID depends upon the PID format used in the fabric:

- When Core PID format is in effect, the area ID for port 0 is 0, for port 1, it is 1 and so forth.

When using Core PID mode on the SilkWorm 12000 (two logical 64-port switches), SilkWorm 48000 (one domain only), and the SilkWorm 24000 configured with two domains, the area IDs for both logical switches (domains) range from 0 to 63. This means that both logical switch 0 and logical switch 1 have a port that is referenced with area ID 0.

For 32-port blades in the SilkWorm 48000 (using FC4-32), the numbering is contiguous up to port 15; from port 16, the numbering is still contiguous, but you must add 128 to each port number. For example, port 16 in slot 1 has a port number and area ID of 128; port number 15 has a port number and area ID of 15.
- When Extended-edge PID format is in effect, the area ID is the port number plus 16 for ports 0 to 111. For port numbers higher than 111, the area ID wraps around so that port 112 has an area ID of 0, and so on. Each 64-port logical switch (domain) has area IDs ranging from 16 to 79. Port numbers higher than 128 are mapped directly to the core PID.

For details about port area IDs in Extended-edge PID mode, refer to [“Changing to Extended Edge PID Format” on page A-10](#).

- If you perform a port swap operation, the port number and area ID no longer match.

To determine the area ID of a particular port, enter the **switchShow** command. This command displays all ports on the current (logical) switch and their corresponding area IDs.

Basic Blade Management

The following sections provide procedures for powering a port blade off and on and for disabling and enabling a port blade.

Powering Port Blades Off and On

Port blades are powered on by default.

To power off a port blade

1. Connect to the switch and log in as admin.
2. Enter the **slotPowerOff** command with the slot number of the port blade you want to power off.

The slot must exist in the logical switch where you are logged in.

Example

```
switch:admin> slotpoweroff 3
Slot 3 is being powered off
switch:admin>
```

To provide power to a port blade

1. Connect to the switch and log in as admin.
2. Enter the **slotPowerOn** command with the slot number of the port blade you want to power on.

The slot must exist in the logical switch where you are logged in.

Example

```
switch:admin> slotpoweron 3
Powering on slot 3
switch:admin>
```

Disabling and Enabling Port Blades

Port blades are enabled by default.

You might need to disable a port blade to perform diagnostics. When diagnostics are executed manually (from the Fabric OS command line), many commands require the port blade to be disabled. This ensures that diagnostic activity does not interfere with normal fabric traffic.

To disable a port blade

1. Connect to the switch and log in as admin.
2. Enter the **slotOff** command with the slot number of the port blade you want to disable:

```
switch:admin> slotoff 3
Slot 3 is being disabled
switch:admin>
```

To enable a port blade

1. Connect to the switch and log in as admin.
2. Enter the **slotOn** command with the slot number of the port blade you want to enable:

```
switch:admin> sloton 3
Slot 3 is being enabled
switch:admin>
```

Conserving Power

To conserve power and ensure that more critical components are the least affected by a power fluctuation, you can power off components in a specified order, using the **powerOffListSet** command.

The available power is compared to the power demand to determine if there is enough power to operate. If there is less power available than the demand, the power-off list is processed until there is enough power for operation. By default, the processing proceeds from slot 1 to the last slot in the chassis. As power becomes available, slots are powered up in the reverse order.



Note

Some FRUs in the chassis may use significant power, yet cannot be powered off through software. For example, a missing blower FRU may change the power computation enough to affect how many slots can be powered up.

The **powerOffListShow** command displays the power off order.

Blade Terminology and Compatibility

Before configuring a chassis, familiarize yourself with the director CP blade and port blade nomenclature, as well as the port blade compatibilities. Often in procedures, only the abbreviated names for CP and port blades are used (for example, the FC4-16 blade). [Table 5-1](#) includes CP and port blade abbreviations and descriptions.

Table 5-1 SilkWorm Director Terminology and Abbreviations

Term	Abbreviation	Blade ID	Definition
SilkWorm 12000 control processor blade	CP1	1	The first generation CP blade provided with the SilkWorm 12000. This CP supports 1 and 2 Gbit/sec port speeds. It only supports the dual domain configuration within the chassis.
SilkWorm 24000 control processor blade	CP2	5	The second generation CP blade provided with the SilkWorm 24000. This CP supports 1 and 2 Gbit/sec port speeds. It supports both the dual domain and a single domain configuration within the chassis.
SilkWorm 48000 control processor blade	CP4	16	The second generation CP blade provided with the SilkWorm 48000. This CP supports 1, 2, and 4 Gbit/sec port speeds, as well as 16 and 32-port blades.
16-port 2-Gbit/sec port blade	FC-16	2	The first generation SilkWorm director 16-port blade supporting 2 Gbit/sec port speeds. This port blade is only compatible with the SilkWorm 12000 or SilkWorm 24000 CP blades.
16-port 2-Gbit/sec port blade	FC2-16	4	The second generation SilkWorm director 16 port blade supporting 1 and 2 Gbit/sec port speeds. This port blade is only compatible with the SilkWorm 24000 or SilkWorm 48000 CP blades.
16-port 4-Gbit/sec port blade	FC4-16	17	The third generation SilkWorm director 16 port blade supporting 1, 2, and 4 Gbit/sec port speeds. This port blade is only compatible with the SilkWorm 24000 or SilkWorm 48000 CP blades. FC4-16 blades do not support private devices.
32-port 4-Gbit/sec port blade	FC4-32	18	A 32 port SilkWorm director port blade supporting 1, 2, and 4 Gbit/sec port speeds. This port blade is only compatible with the SilkWorm 48000 CP blades. FC4-32 blades do not support private devices.
D1 Chassis	n/a	n/a	The first generation chassis. These chassis have a manufacture date prior to January 1, 2004. In Fabric OS 5.0.1, use the chassisShow command to view the backplane revision number for this chassis, 0x1F.
D2 Chassis	n/a	n/a	The second generation chassis. These chassis have a manufacture date from to January 1, 2004 to May 1, 2005. In Fabric OS 5.0, use the chassisShow command to view the backplane revision number for this chassis, 0x1D.
D3 Chassis	n/a	n/a	The third generation chassis. These chassis have a manufacture date from to May 1, 2005 to the present. In Fabric OS 5.0, use the chassisShow command to view the backplane revision number for this chassis, 0x1B.

CP Blades

CP blades determine the director type:

- If CP1 blades are installed, the director is a SilkWorm 12000.
- If CP2 blades are installed, the director is a SilkWorm 24000.
- If CP4 blades are installed, the director is a SilkWorm 48000.

Mixed CP blades are not supported on a single chassis, except during specific upgrade procedures detailed in *SilkWorm Director Migration Guide*. Brocade recommends that each SilkWorm director have only one type of CP blade installed and that each CP (primary and secondary partition) maintains the same firmware version.

Port Blade Compatibility

[Table 5-2](#) indicates which blades are supported for each SilkWorm director.

Table 5-2 Blades Supported by Each SilkWorm Director

Director	Port Blades			
	FC-16	FC2-16	FC4-16	FC4-32
SilkWorm 12000 (CP1)	Supported	NA	NA	NA
SilkWorm 24000 (CP2)	Supported	Supported	Supported	NA
SilkWorm 48000 (CP4)	NA	Supported	Supported	Supported

Setting Chassis Configurations

The **chassisConfig** command allows you to set the chassis configuration for products that support both single-switch (one domain) and dual-switch (two domains) operation.

Table 5-3 lists the supported configurations options for Fabric OS v5.0.1.

Table 5-3 Supported Configuration Options

Option	Number of Domains	Maximum Number of Ports Per Switch	Supported Port Blades	Supported CP Blades	Notes
1	1	128	FC2-16, FC4-16	CP2 or CP4	CP4 fits all chassis except the D2 chassis. Option 1 is the default configuration for SilkWorm 24000.
2	2	64/64	FC2-16	CP2	
3	2	64/64	Left side: FC2-16 Right side: FC-16	CP2	
4	2	64/64	Left side: FC-16 Right side: FC2-16	CP2	
5	1	256	FC4-16, FC4-32	CP4	CP4 fits all chassis except the D2 chassis. Option 5 is the default configuration option for SilkWorm 48000.

Option Result

- 1 One 128-port switch (Blade IDs 4, 17 on slots 1-4, 7-10. Blade ID 5 and 16 on slots 5, 6)
- 2 Two 64-port switches (Blade ID 4 on slots 1-4, 7-10. Blade ID 5 on slots 5, 6)
- 3 Two 64-port switches (Blade ID 4 on slots 1-4; ID 2 on slots 7-10. Blade ID 5 on slots 5, 6)
- 4 Two 64-port switches (Blade ID 2 on slots 1-4; ID 4 on slots 7-10. Blade ID 5 on slots 5, 6)
- 5 One 256-port switch (Blade IDs 4, 17, and 18 on slots 1-4, 7-10. Blade ID 16 on slots 5, 6)

The following sections contain procedures for obtaining chassis information and for configuring director domains using the **chassisConfig** command.

Obtaining Slot Information

For a SilkWorm 12000 or 24000 configured as two logical switches, the chassis-wide commands display or control both logical switches. In the default configuration, SilkWorm 24000 and 48000 directors are configured as one logical switch, so the chassis-wide commands display and control the single logical switch.

To display the status of all slots in the chassis

1. Connect to the switch and log in as user or admin.
2. Enter the **slotShow** command to display the current status of each slot in the system. The format of the display includes a header and four fields for each slot. The fields and their possible values are:

Field	Value
Slot	Displays the physical slot number.
Blade Type	Displays the blade type: SW BLADE: The blade is a switch. CP BLADE: The blade is a control processor. UNKNOWN: The blade is not present or its type is not recognized.
ID	Displays the hardware ID of the blade type.
Status	Displays the status of the blade: VACANT: The slot is empty. INSERTED, NOT POWERED ON: The blade is present in the slot but is turned off. DIAG RUNNING POST1: The blade is present, powered on, and running the post-initialization power-on self test. DIAG RUNNING POST2: The blade is present, powered on, and running the POST (power-on self test). ENABLED: The blade is on and enabled. ENABLED (User Ports Disabled): The blade is on, but external ports have been disabled with the bladeDisable command. DISABLED: The blade is powered on but disabled. FAULTY: The blade is faulty because an error was detected. The reason code numbers displayed are for debugging purposes. UNKNOWN: The blade is inserted but its state cannot be determined.

Configuring a New SilkWorm 24000 with Two Domains

By factory default, the SilkWorm 24000 director is configured as one 128-port switch (one domain). The procedure assumes that the new director:

- Has been installed and connected to power, but is not yet attached to the fabric.
- Has been given an IP address, but is otherwise running factory defaults.

If this is not the case, back up the current configuration before starting, so that you can restore it later if necessary.

- Is running Fabric OS v4.4.0 or later.
- Is running in configuration option one (one switch, FC2-16 cards installed).

Use the following procedure to add a factory new SilkWorm 24000 to a fabric and configure it as two 64-port switches (two domains).

1. Connect to the switch and log in as admin.
2. Enter the **chassisconfig** command without options to verify that the switch is configured with one domain. For example:

```
chassisconfig
Current Option: 1
```

3. Enter the **chassisconfig** command to configure two domains. Use the **-f** option to suppress prompting for uploading the configuration. This command reboots the system.

```
chassisconfig -f 2
Current Option changed to 2
Restoring switch 0 configuration to factory defaults...
All account passwords have been successfully set to factory default.
Restoring switch 1 configuration to factory defaults...
All account passwords have been successfully set to factory default.
```

4. After the system reboots, log in again to the first logical switch (sw0) as admin.
5. Use the **configure** command to configure the sw0 to match your fabric specifications.
If the director is to be merged into an existing fabric, do not configure zoning parameters; these will be propagated automatically when you merge the director into the fabric.
6. Log in to the second logical switch (sw1) as admin.
7. Use the **configure** command to configure the sw1 to match your fabric specifications.
If the director is to be merged into an existing fabric, do not configure zoning parameters; these will be propagated automatically when you merge the director into the fabric.
8. If the fabric is in secure mode, perform the following steps; otherwise, proceed to [step 9](#).
(Refer to the *Secure Fabric OS Administrator's Guide* for specific instructions.)
 - a. Optionally, to configure sw0 and sw1 in one operation, connect them with an ISL link to form a temporary fabric.
 - b. If you want sw0 and sw1 to be fabric configuration servers, update the overall fabric's FCS policy to include them. If not, skip this step.

- c. On sw0, enable security mode and use the **secModeEnable** command to create an FCS list that matches your overall fabric's FCS policy.
 - d. Reset the version stamp on sw0.
 - e. If you connected sw0 and sw1 in [step a](#) and you do not want them connected, disconnect the ISL link between them. If you did not connect them, repeat [step b](#) through [step d](#) on sw1.
9. Optionally, connect the new two-domain SilkWorm 24000 director to the fabric.
 10. Enter the **fabricShow** command to verify that sw0 and sw1 have been merged with the fabric.
 11. Enter the **cfgShow** command to verify that zoning parameters were propagated.

Converting an Installed SilkWorm 24000 to Support Two Domains

Fabric OS versions earlier than v4.4.0 supported only one domain for SilkWorm 24000 directors (one 128-port logical switch). When you upgrade a SilkWorm 24000 director to Fabric OS v4.4.0 or later, you can use the **chassisConfig** command to specify two domains for the director (two 64-port logical switches, sw0 and sw1). This conversion is for SilkWorm 24000 directors using configuration option one (one switch, FC2-16 cards installed).



Note

This procedure restores most configuration parameters to factory defaults. After performing this procedure, you must check the new configuration and reconfigure those parameters that you customized in the old configuration.

During this procedure, power is reset and the CP blades are rebooted, so traffic on the fabric is disrupted. If the fabric is in secure mode, enabling security on the new domains is a complicated task. You should avoid converting existing core switches.

1. Connect to the switch and log in as admin.
2. If the director is already in a fabric, minimize disruption by removing the director from the fabric using one of the following methods:
 - Physically disconnect the director.
 - Use the **portCfgPersistentDisable** command on all connected remote switches to persistently disable their ports that are connected to the director, or remove ISLs that connect the SilkWorm 24000 to the current fabric.
3. Enter the **chassisConfig** command to change the configuration from the default (one domain) to two domains. This command reboots the system.

```
chassisconfig 2
```

During the conversion, you are prompted to save the configuration of sw0. Follow the prompts to save the configuration file.

4. Enter the **ipAddrSet** command to set and confirm the IP address of sw1 (sw1 takes on a default that must be corrected). The IP address of sw0 will already be set.

5. After the system reboots, log in again as admin to each logical switch and enter the **switchName** command to assign a name to the new switch.
6. Using the configuration file saved in [step 3](#) as a guide, manually reconfigure sw0 and sw1.
Do not configure zoning parameters; these are propagated automatically when you merge the director into the fabric.
7. If the fabric is in secure mode, perform the following steps; otherwise, proceed to [step 8](#).
 - a. Optionally, to configure sw0 and sw1 in one operation, connect them with an ISL link to form a temporary fabric.
 - b. If you want sw0 and sw1 to be fabric configuration servers, update the overall fabric's FCS policy to include them. If not, skip this step.
 - c. On sw0, enable security mode and use the **secModeEnable** command to create an FCS list that matches your overall fabric's FCS policy.
 - d. Reset the version stamp on sw0.
 - e. If you connected sw0 and sw1 in [step a](#) and you do not want them connected, disconnect the ISL link between them. If you did not connect them, repeat [step b](#) through [step d](#) on sw1.
8. If you physically disconnected the switch in [step 2](#), reconnect it to the fabric.
If you used the **portCfgPersistentDisable** command in [step 2](#), use the **portCfgPersistentEnable** command to persistently enable all ports that connect the switch to other switches in the fabric.
9. Use the **fabricShow** command to verify that sw0 and sw1 have been merged with the fabric.
10. Use the **configShow** command to verify that zoning parameters were propagated.

Combining SilkWorm 12000 and 24000 Port Blades in One Chassis

For details on combining SilkWorm 12000 and 24000 port blades in a single chassis, refer to the *SilkWorm Director Migration Guide*.

Setting the Blade Beacon Mode

When beaconing mode is enabled, the port LEDs will flash amber in a running pattern from port 0 through port 15 and back again. The pattern continues until the user turns it off. This can be used to locate a particular blade.

To set the blade beacon mode on:

1. Connect to the switch and log in as admin.
2. Enter the **bladeBeacon** command:

```
bladebeacon slotnumber, mode
```

The *slotnumber* is the blade on which you want to enable beacon mode; this slot number must exist on the logical switch. The value 1 turns beaconing mode on, and 0 turns beaconing mode off.

Example

```
switch:admin> bladebeacon 3, 1  
switch:admin>
```

Routing Traffic

This chapter describes SilkWorm switch routing features and procedures.

About Data Routing and Routing Policies

Data moves through a fabric from switch to switch and storage to server along one or more paths that make up a *route*. Routing policies determine the correct path for each frame of data.



Caution

For most configurations, the default routing policy is optimal, and provides the best performance. You should only change the policy if there is a performance issue that is of concern, or a particular fabric configuration requires it.

The following routing policies are available to tune routing performance:

- *Exchange-based* routing

The choice of routing path is based on the SID, DID, and Fibre Channel originator exchange ID (OXID), optimizing path utilization for the best performance. Thus, every exchange can take a different path through the fabric

- *Device-based* routing

The choice of routing path is based on the Fibre Channel addresses of the source device (SID) and the destination device (DID), improving path utilization for better performance. Thus, the same route is always used and the sequence of exchanges is guaranteed.

- *Port-based* routing

The choice of routing path is based only on the incoming port and the destination domain. To optimize port-based routing, the Dynamic Load Sharing feature (DLS) can be enabled to balance the load across the available output ports within a domain.

Device-based and exchange-based routing require the use of DLS; when these policies are in effect, you cannot disable the DLS feature.

Using port-based routing, you can assign a *static route*, in which the path chosen for traffic never changes. In contrast, device-based and exchange-based routing policies always employ *dynamic path selection*. Port-based routing is supported by all SilkWorm models (except SilkWorm 48000 using configuration option 5).

Specifying the Routing Policy

The following routing policies are supported:

- Port-based path selection

Default on SilkWorm 3016, 3250, 3850, 3900, 12000, 24000, 48000 (using configuration options 1 through 4). These switches support the port-based policy only; you cannot change the routing policy for these switches. SilkWorm 200E, 4012, and 4100 switches can also use port-based routing.

- Device-based path selection

SilkWorm 200E, 4012, 4100, and 48000 (using configuration option 5).

If there are devices in your fabric that cannot accommodate out-of-order exchanges, then you should use the device-based policy. Brocade recommends that you also use device-based routing for FICON environments.

- Exchange-based path selection

Default on SilkWorm 200E, 4012, 4100, and 48000 (using configuration option 5)

Refer to [Chapter 5, “Configuring SilkWorm 12000, 24000, and 48000 Directors”](#) for more details about SilkWorm 48000 configuration options.

You can use the **aptpolicy** command to display and specify a different routing policy. Note that if you attempt to set the policy when the SilkWorm 48000 uses configuration options 1-4, an error message is returned. Refer to the *Fabric OS Command Reference Manual* for more details on the **aptpolicy** command.

You must disable the switch before changing the routing policy, and reenabling it afterward.

In the following example, the routing policy for a SilkWorm 4100 is changed from exchange-based to device-based:

```
switch:admin> aptpolicy
Current Policy: 3

3: Default Policy
1: Port Based Routing Policy
2: Device Based Routing Policy
3: Exchange Based Routing Policy
switch:admin> switchdisable
switch:admin> aptpolicy 2
Policy updated successfully.
switch:admin> switchenable
switch:admin> aptpolicy
Current Policy: 2
```

Assigning a Static Route

A static route can be assigned only when the active routing policy is port-based. When device-based or exchange-based routing is active, you cannot assign static routes. Thus, the SilkWorm 48000 using configuration option 5 does not support static routing.

To assign a static route, use the **uRouteConfig** command. To remove a static route, use the **uRouteRemove** command



Note

SilkWorm 3900, 12000, 24000, and 48000 (using configuration options 1 through 4):

When you enter the **uRouteConfig** command, two similar warning messages might display if a platform conflict occurs. The first message displays when the static routing feature detects the conflict. The second message displays when the Dynamic Load Sharing feature detects the condition as it tries to rebalance the route.

A platform conflict occurs if a static route was configured with a destination port that is currently down. The static route is ignored in this case, in favor of a normal dynamic route. When the configured destination port comes back up, the system attempts to reestablish the static route, potentially causing a conflict.

Specifying Frame Order Delivery

The order of delivery of frames is maintained within a switch and determined by the routing policy in effect. Following are the frame delivery behaviors for each routing policy.

- Port-based routing
All frames received on an ingress port destined for a destination domain are guaranteed to exit the switch in the same order in which they were received.
- Device-based routing
All frames received on an ingress port between the same two fabric devices are guaranteed to exit the switch in the same order in which they were received. This policy maintains the order of frames across exchanges between the fabric devices as well.
- Exchange-based routing
All frames received on an ingress port for a given exchange are guaranteed to exit the switch in the same order in which they were received. Because different paths are chosen for different exchanges, this policy does not maintain the order of frames across exchanges.

If even one switch in the fabric delivers out-of-order exchanges, then exchanges are delivered to the target out-of-order, regardless of the policy configured on other switches in the fabric.

In a stable fabric, frames are always delivered in order, even when the traffic between switches is shared among multiple paths. However, when topology changes occur in the fabric (for example, if a link goes down), traffic is rerouted around the failure, and some frames could be delivered out of order. Most destination devices tolerate out-of-order delivery, but some do not.

By default, out-of-order frame-based delivery is allowed to minimize the number of frames dropped. You should only force in-order frame delivery across topology changes if the fabric contains destination devices that cannot tolerate occasional out-of-order frame delivery.

To force in-order frame delivery across topology changes

1. Connect to the switch and log in as admin.
2. At the command line, enter the **iodSet** command.



Note

This command can cause a delay in the establishment of a new path when a topology change occurs; use it with care.

To restore out-of-order frame delivery across topology changes

1. Connect to the switch and log in as admin.
2. Enter the **iodReset** command at the command line.

Using Dynamic Load Sharing

The device-based and exchange-based routing policies depend on the Fabric OS Dynamic Load Sharing feature (DLS) for dynamic routing path selection. When these policies are in force, DLS is by default enabled and cannot be disabled.

When the port-based policy is in force, you can enable DLS to optimize routing. When DLS is enabled, it shares traffic among multiple equivalent paths between switches. DLS recomputes load sharing when a switch boots up or each time an E_Port goes offline and online, or a Fx_Port goes offline.

To check and set DLS

1. Connect to the switch and log in as admin.
2. Enter the **dlsShow** command to view the current DLS setting.

One of the following messages appears:

- “DLS is set” indicates that dynamic load sharing is turned on.
- “DLS is not set” indicates that dynamic load sharing is turned off.

3. Enter the **dlsSet** command to enable DLS or enter the **dlsReset** command to disable it.

You cannot disable DLS when device-based or exchange-based routing policies are in effect.

Example

```
switch:admin> dlsshow
DLS is not set
switch:admin> dlsset
switch:admin> dlsshow
DLS is set
switch:admin> dlsreset
switch:admin> dlsshow
DLS is not set
```

Viewing Routing Path Information

The **topologyShow** and **uRouteShow** commands provide information about the routing path.

1. Connect to the switch and log in as admin.
2. Enter the **topologyShow** command to display the fabric topology, as it appears to the local switch.

The following entries appear:

Local Domain ID	Domain number of the local switch.
Domain	Domain number of the destination switch.
Metric	Cost of reaching the destination domain.
Name	The name of the destination switch.
Path Count	The number of currently active paths to the destination domain.
Hops	The maximum number of hops to reach the destination domain.
Out Port	The Port to which the incoming frame will be forwarded in order to reach the destination domain.
In Ports	Input ports that use the corresponding Out Port to reach the destination domain.
Total Bandwidth	The maximum bandwidth of the out port.
Bandwidth Demand	The maximum bandwidth demand of the in ports.
Flags	Always “D”, indicating a dynamic path.

Example

```

switch:admin> topologyshow
2 domains in the fabric; Local Domain ID: 1
Domain: 6
Metric: 500
Name: switch
Path Count: 4
Hops: 1
Out Port: 60
In Ports: None
Total Bandwidth: 2 Gbps
Bandwidth Demand: 0%
Flags: D
Hops: 1
Out Port: 61
In Ports: None
Total Bandwidth: 2 Gbps
Bandwidth Demand: 0%
Flags: D
Hops: 1
Out Port: 62
In Ports: None
Total Bandwidth: 2 Gbps
Bandwidth Demand: 0%
Flags: D
Hops: 1
Out Port: 58
In Ports: None
Total Bandwidth: 2 Gbps
Bandwidth Demand: 0%
Flags: D

```

3. Use the **uRouteShow** command to display unicast routing information.

SilkWorm 200E, 3016, 3250, 3850, 3900, 4012, and 4100: Use the following syntax:

```
urouteshow [portnumber][, domainnumber]
```

SilkWorm 12000, 24000, and 48000: Use the following syntax:

```
urouteshow [slot/][portnumber][, domainnumber]
```

The following entries appear:

Local Domain	Domain number of the local switch.
In Ports	Port from which a frame is received.
Domain	Destination domain of the incoming frame.
Out Port	The port to which the incoming frame will be forwarded in order to reach the destination domain.
Metric	Cost of reaching the destination domain.
Hops	The maximum number of hops to reach the destination domain.
Flags	Indicates if the route is dynamic (D) or static (S). A static route is assigned using the command uRouteConfig .

Next (Dom, Port) Domain number and port number of the next hop.

This example displays the routing information of all the active ports:

```
switch:admin> urouteshow
Local Domain ID: 3
In PortDomain Out Port Metric Hops Flags Next (Dom, Port)
-----
0 1 11 1000 1 D 1,0
11 2 0 1500 2 D 4,0
4 16 500 1 D 4,0
16 1 27 1000 1 D 1,1
27 2 16 1500 2 D 4,16
4 0 29 500 1 D 4,0
```

This example displays the routing information for port 11 on slot 1:

```
switch:admin> urouteshow 1/11
Local Domain ID: 3
In PortDomain Out Port Metric Hops Flags Next (Dom, Port)
-----
11 2 0 1500 2 D 4,0
4 16 500 1 D 4,0
```

This example displays the routing information of port 11 to domain 4 only:

```
switch:admin> urouteshow 1/11, 4
Local Domain ID: 3
In PortDomain Out Port Metric Hops Flags Next (Dom, Port)
-----
11 4 16 500 1 D 4,0
```

Viewing Routing Information Along a Path

You can display detailed routing information from a source port (or area) on the local switch to a destination port (or area) on another switch. This routing information describes the full path that a data stream travels between these ports, including all intermediate switches.

1. Connect to the switch and log in as admin.
2. Enter the **pathInfo** command. In interactive mode, you can specify the following parameters for display:

Max hops	The maximum number of hops that the pathinfo frame is allowed to traverse.
Domain	The destination domain ID.
Source Port	The port number (or area number for SilkWorm 12000, 24000, or 48000 directors) on which the switch receives frames.
Destination Port	The output port that the frames use to reach the next hop on this path. For the last hop, the destination port.
Basic stats	Basic statistics on every link.

Extended stats	Detailed statistics on every link.
Trace reverse path	Traverses from the destination switch back to the source switches.
Source route	Forces the frame to follow a specified path to reach the destination.
Timeout	The maximum time to wait for a response from pathInfo , in seconds.

Paths always originate on the local switch. The path destination can be specified by domain or port. By default, the path will be the path taken by traffic from the source to destination port, but you can also specify all or portions of a path.

Refer to the *Fabric OS Command Reference Manual* for details on the **pathInfo** command.

This example is from a SilkWorm 3900 switch (other models provide similar information):

```
switch:admin> pathinfo

Max hops: (1..127) [25]
Domain: (1..239) [-1] 1
Source port: (0..255) [-1]
Destination port: (0..255) [-1]
Basic stats (yes, y, no, n): [no]
Extended stats (yes, y, no, n): [no]
Trace reverse path (yes, y, no, n): [no]
Source route (yes, y, no, n): [no]
Timeout: (1..30) [10]

Target port is Embedded

Hop  In Port  Domain ID (Name)      Out Port  BW    Cost
-----
0      E        10 (SW3900)          15        2G    500
1      7         1 (swd3900TechPu)    E         -      -
switch.admin>
```

The information that **pathInfo** provides is:

Hop	The hop number. The local switch is hop 0.
In Port	The port that the frames come in from on this path. For hop 0, the source port.
Domain ID	The domain ID of the switch.
Name	The name of the switch.
Out Port	The output port that the frames use to reach the next hop on this path. For the last hop, the destination port.
BW	The bandwidth of the output ISL, in Gbit/sec. It does not apply to the embedded port.
Cost	The cost of the ISL used by FSPF routing protocol. It only applies to an E_Port.

Administering FICON Fabrics

FICON Overview

IBM Fibre Connections (FICON[®]) is an industry-standard, high-speed input/output (I/O) interface for mainframe connections to storage devices. Fabric OS supports *intermix mode* operations, in which FICON and Fibre Channel technology work together. For specific information about intermix mode and other aspects of FICON, refer to the IBM Redbook, *FICON[®] Native Implementation and Reference Guide*.

Fabric OS provides standard support for FICON single-switch operation.

Multiple-switch cascaded FICON operation requires a Brocade Secure Fabric OS license.

Control Unit Port (CUP) operation requires a Brocade FICON CUP license.

The following Fabric OS *standard features* support FICON fabrics:

- Port swapping
Redirects resources from a failed port to a healthy port without changing the FICON host configuration. Port swapping is available for both FICON and open system environments. Port swapping resolves situations in which the hardware has failed and the channel configurations cannot be changed quickly. Port swapping has minimal or no impact on other switch features.
- Insistent domain ID (IDID)
Allows the switch to insist on a specific domain ID before joining a fabric. This feature guarantees that a switch operates only with its preassigned domain ID.
- The FICON MIB module
Addresses link incident data for FICON hosts and devices connected to a switch. It supplements other MIBs used to manage switches and should be used in conjunction with those other MIBs. For more information, refer to the *Fabric OS MIB Reference Manual*.
- Link incident detection, registration, and reporting
Provide administrative and diagnostic information.

These *optional features* provide further support:

- The Secure Fabric OS optional license
Includes fabric binding, switch binding, and port binding security methods that prevent unauthorized devices from joining a fabric.
- A Fabric Manager optional license
Can be used to manage a fabric that supports FICON and FCP devices and traffic. This is the recommended GUI management tool for FICON environments.
- Web Tools
Can be used to manage a director (switch) that supports FICON and Fibre Channel Protocol (FCP) devices and traffic.

To incorporate and manage FICON on a switch or fabric, your system must have Fabric OS v4.1.2 or later installed. If you are implementing FICON in a single-switch *noncascaded* environment, there are no additional software requirements. The Secure Fabric OS and Advanced Zoning optional Brocade licensed features are required on all switches participating in a FICON multiple-switch *cascaded* environment.



Note

Some licenses are installed and activated on the switch at the factory. Use a Brocade management interface to verify that the required licenses are installed and activated on the switch.

The optional Secure Fabric OS license provides these fabric, switch, and port binding features:

- *Fabric binding* is a security method for restricting switches within a multiple-switch fabric. The Switch Connection Control (SCC) policy prevents unauthorized switches from joining a fabric. Switches are authenticated using digital certificates and unique private keys provided to the Switch Link Authentication Protocol (SLAP).
- *Switch binding* is a security method for restricting devices that connect to a particular switch. If the device is another switch, this is handled by the SCC policy. If the device is a host or storage device, the Device Connection Control (DCC) policy binds those devices to a particular switch. Policies range from completely restrictive to reasonably flexible, based upon customer needs.
- *Port binding* is a security method for restricting host or storage devices that connect to particular switch ports. The DCC policy also binds device ports to switch ports. Policies range from completely restrictive to reasonably flexible, based upon customer needs.

FICON protocol is supported on the following SilkWorm models and Fabric OS releases:

- SilkWorm 3900, Fabric OS v4.1.2 or later.
- SilkWorm 12000, Fabric OS v4.1.2 or later.
- SilkWorm 24000, Fabric OS v4.2.0 or later. The default one-domain configuration is supported; dual domain configurations and SilkWorm 12000 and SilkWorm 24000 port blade configurations are not.

There are two types of FICON configurations:

- A *single-switch* configuration (called *switched point-to-point*) requires that the channel be configured to use single-byte addressing. If the channel is set up for two-byte addressing, then the cascaded configuration setup applies. This type of configuration is described in [“Configuring a Single Switch” on page 7-5](#).
- A *cascaded configuration* (known as a *high integrity fabric*) requires a list of authorized switches. This authorization feature (called *fabric binding*) is available through Brocade Secure Fabric OS. The fabric binding policy allows a predefined list of switches (domains) to exist in the fabric and prevents other switches from joining the fabric. This type of configuration is described in [“Configuring a High-Integrity Fabric” on page 7-5](#).

Control Unit Port (CUP) protocol is used by IBM mainframe management programs to provide in-band management for FICON switches. When it is enabled, you can set up directors in a FICON environment to be managed through IBM mainframe management programs. CUP is an optionally licensed feature available with Fabric OS v4.4.0 or later.

CUP is supported on SilkWorm 3900, 12000, and 24000 models running Fabric OS 4.4.0 or later.

Table 7-1 summarizes the Fabric OS CLI commands that can be used for managing FICON fabrics. For detailed information on these commands, refer to the *Fabric OS Command Reference Manual*.

Table 7-1 Fabric OS Commands Related to FICON and FICON CUP

Command	Description
Standard Fabric OS commands:	
configure	Sets the domain ID and the insistent domain ID mode.
portSwap	Swaps ports.
portSwapDisable	Disables the portSwap command.
portSwapEnable	Enables the portSwap command.
portSwapShow	Displays information about swapped ports.
Commands specific to FICON:	
ficonclear rlir	Removes all RLIR records from the local RLIR database.
ficonclear rnid	Removes all outdated RNID records from the local RNID database.
ficonshow ilir [fabric]	Displays FRU failure information on the local switch or on the fabric.
ficonshow lirr [fabric]	Displays registered listeners for link incidents for the local switch or for the fabric.
ficonshow rlir [fabric]	Displays link incidents for the local switch or for the fabric.
ficonshow rnid [fabric]	Displays node identification data for all devices registered with the local switch or all devices registered with all switches defined in the fabric.
ficonshow switchrnid [fabric]	Displays node identification data for the local switch or for the fabric.
Commands specific to FICON CUP:	
ficoncupset fmsmode	Sets FICON Management Server mode on or off for the switch.
ficoncupset modereg	Sets the mode register bits for the switch.
ficoncupshow fmsmode	Displays the FICON Management Server mode setting for the switch.
ficoncupshow modereg	Displays the mode register bit settings for the switch.



Note

The Fabric OS CLI supports only a subset of the Brocade management features for FICON fabrics. The full set of FICON CUP administrative procedures is available using the Brocade Fabric Manager and Web Tools software features. You can also use an SNMP agent and the FICON Management Information Base (MIB). For information on these tools, refer to:

- Web Tools—*Web Tools Administrator's Guide*
- Fabric Manager—*Fabric Manager User's Guide*
- SNMP Agent and FICON Management Information Base (MIB)—*Fabric OS MIB Reference Manual*

Configuring Switches

This section describes how to configure a switch in a FICON environment. Use the worksheet on [page 7-18](#) to record your configuration information.

The following are recommended FICON environment configuration settings:

- Disable dynamic load sharing (**dlsReset** command).
If DLS is enabled, traffic on existing ISL ports might be affected when one or more new ISLs is added between the same two switches. Specifically, adding the new ISL might result in dropped frames as routes are adjusted to take advantage of the bandwidth provided. By disabling DLS, you ensure that there will be no dropped frames.

A similar situation occurs when an ISL port is taken offline and then brought back online. When the ISL port goes offline, the traffic on that port is rerouted to another ISL with a common destination. When the ISL port comes back online and DLS is enabled, the rerouting of traffic back to the ISL port might result in dropped frames. If DLS is not enabled, traffic will not be routed back.
- Configure ports that are connected to 1-Gbit/sec channels for fixed 1-Gbit/sec speed. Otherwise, when using fixed 1-Gbit/sec channels (both G5 and FICON Express), the FICON host might generate erroneous link incidents when the channels are coming online. These link incidents will result in a call home. Other than the generated link incident, the channel will come online and function normally.
- Enable in-order delivery (**iodSet** command).
- Enable VC translation link initialization on Extended Fabrics links, to stabilize them. Refer to [page 11-8](#) for details on this option of the **portCfgLongDistance** command.
- Although there are no specific zoning rules related to FICON environments, it is recommended that you follow standard FCP zoning practices. For management purposes, put FCP devices in one zone and FICON devices in another zone when operating in a mixed environment.

Preparing a Switch

To verify that a switch is ready to be used in a FICON environment, complete the following steps:

1. Connect to the switch and log in as admin.
2. If not in a cascaded environment, proceed to [step 3](#).

If in a FICON cascaded environment, enter the following commands:

- **licenseShow** to verify that required licenses (Secure Fabric OS and Zoning) are activated
 - **secModeShow** to determine if Secure Fabric OS is enabled; if it is disabled, enable it
 - **secPolicyShow** to verify that the SCC_POLICY is active
3. Enter **switchShow** to verify that the switch and devices are online.
 4. Enter **ficonshow rnid** to verify that the FICON devices are registered with the switch.
 5. Enter **ficonshow lirr** to verify that the FICON host channels are registered to listen for link incidents.
 6. Optionally, to use FICON CUP, refer to “[Using FICON CUP](#)” on [page 7-9](#).

Configuring a Single Switch

Single-switch configuration does not require IDID or fabric binding, provided that connected channels are configured for single-byte addressing. However, you should configure IDID to ensure that domain IDs are maintained.

Configuring a High-Integrity Fabric

To configure a high-integrity fabric (cascaded configuration):

1. Disable each switch in the fabric.
2. For each switch:
 - a. Enable the IDID flag.
 - b. Set the domain ID.
 - c. Install security certificates and keys.
3. Enable the switches; this builds the fabric.
4. Set up security on the primary FCS switch.

Use Quickmode for FICON, which activates the SCC_POLICY and does not create a DCC policy. The security policies are distributed to each switch in the fabric. (For details on the Quickmode procedure, refer to the *Secure Fabric OS Administrator's Guide*.)

5. Connect and enable channel and control unit (CU) devices. The Query for Security Attributes (QSA) response to the channel indicates that the fabric binding and IDID are enabled.

[Figure 7-1](#) and [Figure 7-2](#) show two viable cascaded configurations. These configurations require Channel A to be configured for two-byte addressing and require IDID and fabric binding. There can be only two switches in the path from the channel to the control unit.

Figure 7-1 Cascaded Configuration, Two Switches

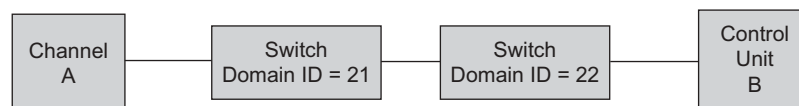
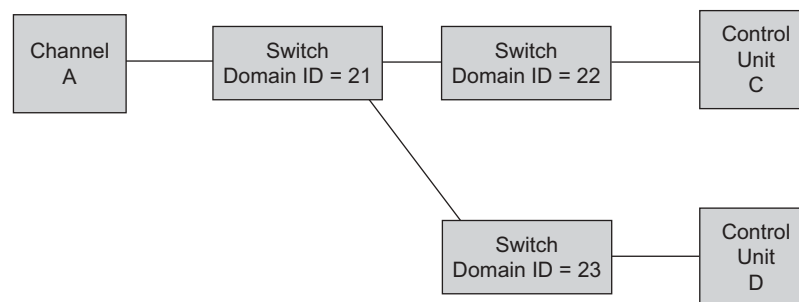


Figure 7-2 Cascaded Configuration, Three Switches



Setting a Unique Domain ID

In a cascaded configuration, each switch must have a unique domain ID, and insistent domain ID (IDID) mode must be enabled. To set a unique domain ID and enable IDID mode, follow these steps:

1. Connect to the switch and log in as admin.
2. Verify that the switch has a unique domain ID. If it does not, set a unique domain ID.

For instructions on displaying and changing the domain ID, refer to [“Working With Domain IDs” on page 2-15](#).

3. Enter the **switchDisable** command to disable the switch.
4. Enter the **configure** command.
5. Enter **y** after the Fabric Parameters prompt.
6. To enable IDID mode, enter **y** after the “Insistent Domain ID Mode” prompt.
(You can disable this mode by entering **n**.)
7. Respond to the remaining prompts (or press **Ctrl-d** to accept the other settings and exit).
8. Enter the **switchEnable** command to reenable the switch.

Example

```
switch:admin> configure
Configure...
Fabric parameters (yes, y, no, n): [no] yes
Domain: (1..239) [3] 5
  R_A_TOV: (4000..120000) [10000]
  E_D_TOV: (1000..5000) [2000]
  Data field size: (256..2112) [2112]
  Sequence Level Switching: (0..1) [0]
  Disable Device Probing: (0..1) [0]
  Suppress Class F Traffic: (0..1) [0]
  VC Encoded Address Mode: (0..1) [0]
  Per-frame Route Priority: (0..1) [0]
  Long Distance Fabric: (0..1) [0]
  BB credit: (1..16) [16]
Insistent Domain ID Mode (yes, y, no, n): [no] y
Virtual Channel parameters (yes, y, no, n): [no]
Switch Operating Mode (yes, y, no, n): [no]
Zoning Operation parameters (yes, y, no, n): [no]
RSCN Transmission Mode (yes, y, no, n): [no]
Arbitrated Loop parameters (yes, y, no, n): [no]
System services (yes, y, no, n): [no]
Portlog events enable (yes, y, no, n): [no]
Committing configuration...done.
switch:admin>
```

Displaying Information

You can display link incidents, registered listeners, node identification data, and FRU failures, as described in the following procedures.

Link Incidents

The registered link incident record (RLIR) ELS contains the link incident information sent to a listener N_Port.

To display link incidents, connect to the switch, log in as user, and enter one of the following commands:

- For the local switch: **ficonshow rlir**
- For all switches defined in the fabric: **ficonshow rlir fabric**

Registered Listeners

To display registered listeners for link incidents, connect to the switch, log in as user, and enter one of the following commands:

- For the local switch: **ficonshow lirr**
- For all switches defined in the fabric: **ficonshow lirr fabric**

Node Identification Data

To display node-identification data, connect to the switch, log in as user, and enter any of the following commands:

- For the local switch: **ficonshow switchrnid**
- For all switches defined in the fabric: **ficonshow switchrnid fabric**
- For all devices registered with the local switch: **ficonshow rnid**
- For all devices registered with all switches defined in the fabric: **ficonshow rnid fabric**

FRU Failures

To display FRU failure information, connect to the switch, log in as admin, and enter one of the following commands:

- For the local switch: **ficonshow ilir**
- For all switches defined in the fabric: **ficonshow ilir fabric**

Swapping Ports

If a port malfunctions, or if you want to connect to different devices without having to re-wire your infrastructure, you can move a port's traffic to another port (*swap ports*) without changing the I/O Configuration Data Set (IOCDS) on the mainframe computer.

To swap ports, perform the following steps (refer to the example that follows):

1. Connect to the switch and log in as admin.
2. Enter the **portSwapEnable** command (to enable the command for port swapping).
3. Enter the **portDisable** command to disable the two ports to be swapped.
4. Enter the **portSwap** command to swap the ports.

Any port in the switch can be used as the alternate for any other port within the same switch.

5. Reenable the ports using the **portEnable** command.
6. Enter **portSwapDisable** (to disable the command for port swapping).

Example

In the following example:

- *slot* is the slot number of the port blade for a system with port blades (optional).
- *portA* is the original port number.
- *portB* is the alternate port number.

```
switch:admin> portswapenable
switch:admin> portdisable [slot/] portA [slot/]portB
switch:admin> portswap [slot/] portA [slot/]portB
switch:admin> portenable [slot/] portA [slot/]portB
switch:admin> portswapdisable
```

You can use the **portSwapShow** command to display information about swapped ports in a switch.

You can use the **portSwap** command to disable the portswap feature. You cannot use the **portSwap** command after this feature is disabled. The enabled state of the portswap feature is persistent across reboots and power cycles. Enabling and disabling the portswap feature does not affect previously executed portswap operations.

Refer to the *Fabric OS Command Reference Manual* for additional details about the **portSwap** command.

Clearing the FICON Management Database

You can clear RLIR and RNID records from the FICON management database as follows:

1. Connect to the switch and log in as admin.
2. To remove all the RLIR records from the local RLIR database, enter **ficonclear rlir**.
3. To remove all the RNID records marked “not current” from the local RNID database, enter **ficonclear rnid**.

Using FICON CUP

Host-based management programs manage switches using CUP protocol by sending commands to an emulated control device in Fabric OS. A Brocade switch that supports CUP can be controlled by one or more host-based management programs, as well as by Brocade tools.

A *mode register* controls the behavior of the switch with respect to CUP itself, and with respect to the behavior of other management interfaces.

FICON Management Server mode (fmsmode) must be enabled on the switch to enable CUP management features. When this mode is enabled, Fabric OS prevents local switch commands from interfering with host-based management commands by initiating *serialized access* to switch parameters.

If more than one switch is to be used in the FICON CUP fabric, Secure Fabric OS must be installed. Refer to [“Configuring a High-Integrity Fabric” on page 7-5](#) for more information.

If Brocade Advanced Zoning is in use, refer to [“Zoning and PDCM Considerations” on page 7-15](#).

Setup Summary

To set up FICON CUP, perform the following actions *in the order stated*:

1. Install Fabric OS v4.4.0 or later on a SilkWorm switch.
2. **For SilkWorm 24000 only:** Use the **portDisable** command to disable (block) port 126.

Port 126 is not supported in a CUP environment. After fmsmode has been successfully enabled, port 126 remains disabled. It cannot be used either as an F_Port or an E_Port. Because port 126 is not available after enabling fmsmode, you should first move any fiber connected to port 126 to another free port.
3. Install a CUP license on the switch. Refer to [“Maintaining Licensed Features” on page 2-8](#).
4. Enable FICON management server mode (fmsmode) on the switch. Refer to [“Enabling and Disabling FICON Management Server Mode” on page 7-9](#).

After completing the setup, you can configure CUP attributes (FMS parameters). Refer to [“Setting Mode Register Bits” on page 7-13](#).

Enabling and Disabling FICON Management Server Mode

To enable fmsmode:

1. Connect to the switch and log in as admin.
2. Enter **ficoncupset fmsmode enable**.

To disable fmsmode:

1. Connect to the switch and log in as admin.
2. Enter **ficoncupset fmsmode disable**.

The `fmsmode` setting can be changed whether the switch is offline or online. If `fmsmode` is changed while the switch is online, a device reset is performed for the control device and an RSCN is generated with PID 0xDDFE00 (where 0xDD is the domain ID of the switch).

When FICON Management Server mode is on, the Fabric OS CLI commands listed here return a “switch busy” response if they are issued when the host-based management tool is performing a write operation. This *serialization* prevents interference from local switch commands when a host-based management program is being used to administer the switch.

<code>bladeDisable</code>	<code>slotOff</code>
<code>bladeEnable</code>	<code>slotOn</code>
<code>portDisable</code>	<code>switchCfgPersistentDisable</code>
<code>portEnable</code>	<code>switchDisable</code>
<code>portName</code>	<code>switchEnable</code>
<code>portShow</code>	<code>switchName</code>
<code>portSwap</code>	<code>switchShow</code>



Note

You cannot use the **`portCfgPersistentEnable`** and **`portCfgPersistentDisable`** commands to persistently enable and disable ports when FICON Management Server mode is on. Refer to the procedure [“Persistently Enabling/Disabling Ports” on page 7-13](#).

Changing `fmsmode` from *disabled* to *enabled* triggers the following events:

- Access to switch parameters is serialized.
- The active CUP configuration data is established as follows:
 - Port and switch names are not read from the IPL; they remain as previously set.
 - Port “Block” and “Unblock” values are not read from the IPL; they remain as previously set with the **`portEnable`** and **`portDisable`** commands.
 - PDCM values are read from the IPL; the default is “Allow All.”
- Brocade Advanced Zoning, if used, continues to be in force. If there are any differences in restrictions set up with Brocade Advanced Zoning and PDCM, the most restrictive rules are automatically applied.
- RSCNs are sent to devices if PDCM results in changes to connectivity between a set of ports.

Changing `fmsmode` from *enabled* to *disabled* triggers the following events:

1. A device reset is performed on the control device.
2. PDCM is no longer enforced.
3. RSCNs might be generated to some devices if PDCM removal results in changes to connectivity between a set of ports.
4. If a given port was set to “Block” or “Unblock,” that port remains disabled or enabled.
5. Serialized access to switch parameters ceases.

Setting Up CUP When FICON Management Server Mode Is Enabled

Fmsmode may be enabled and in use on a switch without a CUP License. The transition from fmsmode disabled to fmsmode enabled with the CUP license installed triggers the notification to the host systems that the CUP feature is available. Without this notification the host systems will never know the CUP feature is available, and consequently will never try to communicate with it. Hence, it is possible that fmsmode may already be enabled on the switch.

If FICON Management Server mode is already enabled, set up CUP as follows:

1. Verify that FICON Management Server mode is enabled by entering the **ficoncupshow fmsmode** command.

If FICON Management Server mode is not enabled, refer to [“Enabling and Disabling FICON Management Server Mode” on page 7-9](#).



Caution

If fmsmode is already enabled, disabling it might be disruptive to operation because ports that were previously prevented from communicating will now be able to do so.

2. If FICON Management Server mode is enabled, then disable it by entering the **ficoncupset fmsmode disable** command.

Install a CUP license key as described in [“Adding and Removing FICON CUP Licenses” on page 7-14](#).

3. Enter the **ficoncupset fmsmode enable** command.

Displaying the fmsmode Setting

The **ficoncupshow fmsmode** command displays the effective fmsmode setting for the switch.

Example

```
switch:admin> ficoncupshow fmsmode
fmsmode for the switch: Enabled
```

Displaying Mode Register Bit Settings

The mode register bits are described in [Table 7-2](#).

Table 7-2 FICON CUP Mode Register Bits

POSC	Programmed offline state control. When this bit is set on, the host is prevented from taking the switch offline. The default setting is 1 (on).
UAM	User alert mode. When this bit is set on, a warning is issued when an action is attempted that will write CUP parameters on the switch. The default setting is 0 (off).
ASM	Active=saved mode. When this bit is set on, all CUP configuration parameters are persistent, meaning that they will be saved in nonvolatile storage in the initial program load (IPL) file that is applied upon a cold reboot or a power cycle. The default setting is 1 (on).
DCAM	Switch clock alert mode. When this bit is set on, a warning is issued when the date , tsClockServer , or tsTimeZone commands are entered to set the time and date on the switch. The default setting is 0 (off).
ACP	Alternate control prohibited. Because the Fabric OS CLI, Web Tools, and Fabric Manager are considered to be switch consoles, this bit has no effect on their operation. Attempts to set CUP parameters through SNMP are denied when this bit is set on. The default setting is 1 (on).
HCP	Host control prohibited. When this bit is set on, the host is not allowed to set CUP parameters. The default setting is 1 (on).

The **ficoncupshow modereg** command displays the mode register bit settings for the switch. A display of 0 indicates that the mode register bit is set off; 1 indicates that the bit is set on.

The command format is:

```
ficoncupshow modereg [bitname]
```

where *bitname* is one of the mode register bits described in [Table 7-2](#).

To display all mode register bit settings for the switch:

```
switch:admin> ficoncupshow modereg
POSC  UAM  ASM  DCAM  ACP  HCP
-----
    1    0    1    0    1    1
```

To display the mode register bit HCP for the switch:

```
switch:admin> ficoncupshow modereg HCP
HCP
1
```

Setting Mode Register Bits

The **ficoncupset modereg** command is used to set the FICON CUP mode register bits for the local switch. Consider the following when changing mode register bits:

- As required by the CUP protocol, the UAM bit cannot be changed using this command.
- All mode register bits except UAM are saved across power on/off cycles; the UAM bit is reset to 0 following a power-on.
- Mode register bits can be changed when the switch is offline or online. If the ACP or HCP bits are changed when the switch is online, they will take effect any time between the completion of the current command and the end of the CCW command chain (or the next alternate manager operation).

The command format is:

```
ficoncupset modereg [bitname] 0 | 1
```

where:

<i>bitname</i>	One of the mode register bits described in Table 7-2 on page 7-12 .
0	Specifies that the bit is off.
1	Specifies that the bit is on.

The following example sets the mode register bit HCP to off:

```
switch:admin> ficoncupset modereg HCP 0
Mode register bit HCP has been set to 0.
```

The following example sets the mode register bit ACP to on:

```
switch:admin> ficoncupset modereg ACP 1
Mode register bit ACP has been set to 1.
```

Persistently Enabling/Disabling Ports

When **fmsmode** is enabled, you cannot use the **portCfgPersistentEnable** and **portCfgPersistentDisable** commands to persistently enable and disable ports. Instead, use this procedure:

1. Enter the following command to display the mode register bit settings:

```
ficoncupshow modereg
```

2. Verify that the ASM bit is set on (1).
3. If the ASM bit is set off (0), enter the following command to set it on:

```
ficoncupset modereg asm 1
```

4. Use the **portEnable** and **portDisable** commands to enable and disable ports as necessary.

The ports remain enabled or disabled after a switch reboot.

In this example, the ASM bit is set to on; then, the port at slot 1, port 1 is enabled persistently:

```
switch:admin> ficoncupshow modereg
POSC  UAM  ASM  DCAM  ACP  HCP
-----
      1   0   0   0   1   1

switch:admin> ficoncupset modereg ASM 1
Mode register bit ASM has been set to 1.

switch:admin> portenable 1/1
```

Port and Switch Naming Standards

Fabric OS handles differences in port and switch naming rules between CUP and itself as follows:

- CUP employs 8-bit characters in port address names and switch names; Fabric OS employs 7-bit characters. When fmsmode is enabled, all characters greater than 0x40 and not equal to 0xFF (EBCDIC code page 37 [0x25]) are allowed in the name; therefore, it is possible for a channel to set a name with nonprintable characters. If a name contains nonprintable characters, they are displayed as dots (...). The following characters are also displayed as dots: semicolon (;), comma (,), “equal” sign (=), and “at” sign (@).



Note

Configuration files that contain nonprintable characters should not be edited manually, because many editors replace nonprintable characters with some other characters without warning the user first.

- CUP has a 24-character unique port name limitation; Fabric OS supports port names up to 32 characters long. When fmsmode is enabled, names longer than 24 characters are truncated.
- To ensure that they are unique, the characters ~00, ~01, ~02, and so on are appended to port names.
- CUP allows a 24-character switch name; Fabric OS limits the switch name to 15 characters. To reconcile this difference, Fabric OS files the first 15 characters in the WWN record and stores the extra characters for CUP use.

Adding and Removing FICON CUP Licenses

If fmsmode is enabled when the FICON CUP license is removed, the control device is reset. PDCM enforcement continues. If fmsmode is disabled when the FICON CUP license is removed, no special action is taken.

If fmsmode is enabled on a switch that does not have a FICON CUP license and then the license is installed, you must first disable and then reenable fmsmode. If fmsmode is disabled and a FICON CUP license is installed, no special action is required.

Zoning and PDCM Considerations

The FICON Prohibit Dynamic Connectivity Mask (PDCM) controls whether or not communication between a pair of ports in the switch is prohibited or allowed. If there are any differences in restrictions set up with Brocade Advanced Zoning and PDCM, the most restrictive rules are automatically applied.

All FICON devices should be configured in a single zone using the “Domain, Area” notation. PDCM can then be used to “Allow” or “Prohibit” access between specific port pairs.

PDCM persists across a failover because it is replicated at all times to the standby CP blade. The active PDCM configuration is saved to the IPL if the ASM bit is set on.

Backing Up and Restoring Configurations

The Fabric OS **configUpload** command saves up to 16 FICON configuration files, including IPL files. For details on the behavior of the **configDownload** command, refer to [“Restoring Configurations in a FICON Environment” on page 4-4](#).

Troubleshooting

The following sources provide useful problem-solving information:

- The standard support commands (**portLogDump**, **supportSave**, **supportShow**) or the Fabric Manager Event Log

By default, the FICON group in the supportShow output is disabled. To enable the capture of FICON data in the supportShow output, enter the **supportshowcfgenable ficon** command. After you get confirmation that the configuration has been updated, the following will be collected and appear in the output for the supportShow command:

- ficoncupshow fmsmode
- ficoncupshow modereg
- ficonDbg dump rnid
- ficonDbg log
- ficonShow ilir
- ficonShow lirr
- ficonShow rlir
- ficonShow rnid
- ficonShow switchrnid
- ficucmd dump -A
- Other detailed information for protocol-specific problems:
 - Display port data structures using the **ptDataShow** command.
 - Display port registers using the **ptRegShow** command.

Identifying Ports

The **ficonshow rlir** command displays, among other information, a tag field for the switch port. You can use this tag to identify the port on which a FICON link incident occurred. The tag field is a concatenation of the switch domain ID and port number, in hexadecimal format. The following example shows a link incident for the switch port at domain ID 120, port 93 (785d in hex):

```
switch:admin> ficonshow rlir
{
  {Fmt  Type PID      Port      Incident Count  TS Format  Time Stamp
   0x18 F    785d00   93              1 Time server Thu Apr 22 09:13:32 2004
   Port Status:          Link not operational
   Link Failure Type:     Loss of signal or synchronization

   Registered Port WWN      Registered Node WWN      Flag  Node Parameters
   50:05:07:64:01:40:16:03  50:05:07:64:00:c1:69:ca  0x10  0x200115
   Type number:             002064
   Model number:            103
   Manufacturer:            IBM
   Plant of Manufacture:    02
   Sequence Number:         0000000169CA
   tag:                     155d

   Switch Port WWN          Switch Node WWN          Flag  Node Parameters
   20:5d:00:60:69:80:45:7c  10:00:00:60:69:80:45:7c  0x00  0x200a5d
   Type number:             SLKWRM
   Model number:            24K
   Manufacturer:            BRD
   Plant of Manufacture:    CA
   Sequence Number:         000000000078
   tag:                     785d
}
}
The Local RLIR database has 1 entry.
switch:admin> ficonshow rlir
```

Backing Up FICON Files

The FICON file access facility is used to store configuration files. This includes IPL and other configuration files. The Fabric OS saves the IPL and all other configuration files on the switch. A maximum of 16 configuration files, including the IPL file, are supported.

You can upload the configuration files saved on the switch to a management workstation using the **configUpload** command. If the switch loses the configuration due to a hardware failure or filesystem error, use the **configDownload** command to restore previously uploaded configuration files. Because data uploaded using the **configUpload** command also contains the IPL, if *Active=Saved* mode is enabled, then the switch will ignore the IPL file downloaded via the **configDownload** command.

To upload the configuration files

Enter the **configUpload** command.

When you execute the **configUpload** command, all the files saved in the file access facility are uploaded to a management workstation (there is a section in the uploaded configuration file labeled “FICON_CUP” that exists in an encoded format).

To download configuration files with Active=Saved mode enabled

Enter the **configDownload** command.

The contents of existing files saved on the switch, which are also present in the “FICON_CUP” section, are overwritten.

The files in the “FICON” section of the configuration file, which are not currently on the switch, are saved on the switch.

The IPL will not be replaced because *Active=Saved* mode is enabled. A warning message is displayed in the event log to warn users that the IPL will not be overwritten.

To download configuration files with Active=Saved mode disabled

Enter the **configDownload** command.

The contents of existing files saved on the switch, which are also present in the “FICON_CUP” section, are overwritten.

The files in the “FICON” section of configuration file, which are not currently on the switch, are saved on the switch.

The IPL will be replaced because *Active=Saved* mode is disabled.

Recording Configuration Information

You can use the following worksheet for recording FICON configuration information.

[illegible]

Sample IOCP Configuration File for SilkWorm 3900, 12000, and 24000 Switches

The channel subsystem controls communication between a configured channel, the control unit, and the device. The I/O Configuration Dataset (IOCDS) defines the channels, control units, and devices to the designated logical partitions (LPARs) within the server; this is defined using the Input/Output Configuration Program (IOCP). The IOCP statements are typically built using the hardware configuration dialog (HCD). The interactive dialog is used to generate your Input/Output Definition File (IODF), invoke the IOCP program, and subsequently build your production IOCDS.

Each FICON director in a fabric must have a unique domain ID and a unique switch ID. The switch ID used in the IOCP definitions can be any value between x'00' to x'FF'. The domain ID range for directors is hex x'01' to x'EF' or decimal 1 to 239. When defining the switch IDs in the IOCP definitions, ensure that you use values within the FICON director's range.

The switch ID has to be assigned by the user and must be unique within the scope of the definitions (IOCP and HCD).

The domain ID is assigned by the manufacturer and can be customized to a different value. It must be unique within the fabric.

Brocade recommends that the switch ID (in IOCP or HCD) be set to the same value as the domain ID of the FICON director, which is defined to the FICON director at installation time. This simplifies the configuration and reduces confusion by having a common designation across all definitions.

For more information, refer to the IBM publication *zSeries Input/Output Configuration Program User's Guide for ICP IOCP* (SB10-7037).

In the following sample IOCP configuration file, the UNIT value for FICON CUP definitions is **2032** for any FICON director regardless of vendor or platform. So all SilkWorm 3900, SilkWorm 12000, or SilkWorm 24000 switches require **UNIT=2032** for the CUP definition. All Domain IDs are specified in hex values in the IOCP (and not in decimal values); the Domain IDs in the example are for demonstration purposes only.

```
*-----
* SilkWorm 24000 Domain_ID=61 (in hex)
*-----
CNTLUNIT CUNUMBR=0D8,UNITADD=00,UNIT=2032,
        PATH=(50,51),
        LINK=(61FE,61FE)
IODEVICE ADDRESS=(0D8,1),CUNUMBR=0D8,UNIT=2032,STADET=Y,UNITADD=00
*-----
*SilkWorm 12000 Domain_ID=22 (left side logical switch 0 in hex)
*-----
CNTLUNIT CUNUMBR=0D9,UNITADD=00,UNIT=2032,
        PATH=(8A,8B),
        LINK=(22FE,22FE)
IODEVICE ADDRESS=(0D9,1),CUNUMBR=0D9,UNIT=2032,STADET=Y,UNITADD=00
*-----
* SilkWorm 12000 Domain_ID=23 (right side logical switch 1 in hex)
*-----
*
CNTLUNIT CUNUMBR=0DA,UNITADD=00,UNIT=2032,
        PATH=(92,93),
        LINK=(23FE,23FE)
IODEVICE ADDRESS=(0DA,1),CUNUMBR=0DA,UNIT=2032,STADET=Y,UNITADD=00
*-----
* SilkWorm 3900 Domain_ID=25 (in hex)
*-----
*
CNTLUNIT CUNUMBR=0DB,UNITADD=00,UNIT=2032,
        PATH=(5A,5B),
        LINK=(25FE,25FE)
IODEVICE ADDRESS=(0DB,1),CUNUMBR=0DB,UNIT=2032,STADET=Y,UNITADD=00
*
*-----
```

Sample RMF Configuration File for Mainframe

Resource management facility (RMF) is a performance management tool that measures selected areas of system activity, including different views of the FICON channel. RMF presents data collected in the form of System Management Facility (SMF) records. This data is essential for any kind of FICON channel performance troubleshooting.

To obtain an RMF FICON director activity report, you must include the keyword “FCD” in the RMF configuration file for the FICON director (this is generic for any FICON director). You must also define the CUP port. In the sample below, the keyword is boldfaced.

```

/*****
/*  MONITOR I OPTIONS                                     */
/*                                                     */
/*              X A      O N L Y                       */
/*                                                     */
/*****

      FCD                      /* FICON Director      */
CHAN                      /* COLLECT CHANNEL STATISTICS */
CPU                      /* COLLECT CPU STATISTICS    */
CYCLE(1000)              /* SAMPLE ONCE EVERY SECOND  */
DEVICE(NOSG)             /* PREVENT SORT OF STORAGE GROUPS*/
DEVICE(NOCHRDR)          /* CHARACTER READER STATISTICS
                        WILL NOT BE COLLECTED          */
DEVICE(COMM)             /* COMMUNICATION EQUIPMENT STATS.
                        WILL BE COLLECTED PDS 1/25/94 */
DEVICE(DASD)             /* DIRECT ACCESS DEVICE STATISTICS
                        WILL BE COLLECTED          */
DEVICE(GRAPH)            /* GRAPHICS DEVICE STATISTICS
                        WILL BE COLLECTED          */
DEVICE(TAPE)             /* TAPE DEVICE STATISTICS
                        WILL BE COLLECTED          */
DEVICE(NOUNITR)          /* UNIT RECORD DEVICE STATISTICS
                        WILL NOT BE COLLECTED        */
DEVICE(NONMBR)           /* NO DEVICE SELECTIVITY BY
                        DEVICE NUMBERS              */
IOQ(DASD)               /* COLLECT DASD I/O QUEUING
                        STATISTICS                  */
IOQ(NOCHRDR)            /* PREVENT CHARACTER READER
                        I/O QUEUING STATISTICS      */
IOQ(NOUNITR)            /* PREVENT UNIT RECORD DEVICE
                        I/O QUEUING STATISTICS      */
IOQ(COMM)               /* COLLECT COMMUNICATION EQUIPMENT
                        I/O QUEUING STATS PDS 1/25/94 */
IOQ(GRAPH)              /* COLLECT GRAPHICS DEVICE

```


Configuring the Distributed Management Server

The Brocade Fabric OS Distributed Management Server allows a SAN management application to retrieve information and administer interconnected switches, servers, and storage devices. The management server assists in the autodiscovery of switch-based fabrics and their associated topologies.

A client of the management server can find basic information about the switches in the fabric and use this information to construct topology relationships. The management server also allows you to obtain certain switch attributes and, in some cases, modify them. For example, logical names identifying switches can be registered with the management server.

The management server provides several advantages for managing a Fibre Channel fabric:

- It is accessed by an external Fibre Channel node at the well-known address *FFFFFAh*, so an application can access information about the entire fabric management with minimal knowledge of the existing configuration.
- It is replicated on every SilkWorm switch within a fabric.
- It provides an unzoned view of the overall fabric configuration. This fabric topology view exposes the internal configuration of a fabric for management purposes; it contains interconnect information about switches and devices connected to the fabric. Under normal circumstances, a device (typically an FCP initiator) queries the Name Server for storage devices within its member zones. Because this limited view is not always sufficient, the management server provides the application with a list of the entire Name Server database.



Note

Management server platform service is available only with Fabric OS v2.3.0 and later.

Enabling and Disabling the Platform Services

The management server is located at the Fibre Channel well-known address *FFFFFAh*. All management services except platform services are enabled by default.

To enable platform services

1. Connect to the switch and log in as admin.
2. Enter the **msplMgmtActivate** command.

```
switch:admin> msplmgmtactivate
Request to activate MS Platform Service in progress.....
*Completed activating MS Platform Service in the fabric!
switch:admin>
```

To disable platform services

1. Connect to the switch and log in as admin.
2. Enter the **msplMgmtActivate** command.
3. Enter **y** to confirm the deactivation.

```
switch:admin> msplmgmtdeactivate
MS Platform Service is currently enabled.
This will erase MS Platform Service configuration
information as well as database in the entire fabric.
Would you like to continue this operation? (yes, y, no, n): [no] y
Request to deactivate MS Platform Service in progress.....
*Completed deactivating MS Platform Service in the fabric!
switch:admin>
```

Controlling Access

You can use the **msConfigure** command to control access to the management server database.

An access control list (ACL) of WWN addresses determines which systems have access to the management server database. The ACL typically contains those WWNs of host systems that are running management applications.

If the list is empty (the default), the management server is accessible to all systems connected in-band to the fabric. For more access security, you can specify WWNs in the ACL so that access to the management server is restricted to only those WWNs listed.

The ACL is switch-based. Therefore, only hosts that are connected directly to the switch are affected by the ACL. A host that is somewhere else in the fabric and is connected to a switch with an empty ACL is allowed to access the management server.



Note

The **msConfigure** command is disabled if the switch is in secure mode. Refer to the *Secure Fabric OS Administrator's Guide* for more information.

To display the management server ACL

1. Connect to the switch and log in as admin.
2. Enter the **msConfigure** command.
The command becomes interactive.
3. At the select prompt, enter **1** to display the access list.

A list of WWNs that have access to the management server is displayed.

In the following example, the list is empty:

```
switch:admin> msconfigure
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [1] 1
MS Access list is empty.
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
done ...
switch:admin>
```

To add a member to the ACL

1. Connect to the switch and log in as admin.
2. Enter the **msConfigure** command.
The command becomes interactive.
3. At the select prompt, enter **2** to add a member based on its port/node WWN.
4. Enter the WWN of the host to be added to the ACL.
5. At the prompt, enter **1** to verify the WWN you entered was added to the ACL.
6. After verifying that the WWN was added correctly, enter **0** at the prompt to end the session.
7. At the “Update the FLASH?” prompt, enter **y**.
8. Press **Enter** to update the nonvolatile memory and end the session.

Example

```

switch:admin> msconfigure
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [1] 2
Port/Node WWN (in hex): [00:00:00:00:00:00:00:00] 20:00:00:20:37:65:ce:aa
*WWN is successfully added to the MS ACL.
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [2] 1
MS Access List consists of (14): {
  20:00:00:20:37:65:ce:aa
  20:00:00:20:37:65:ce:bb
  20:00:00:20:37:65:ce:ff
  20:00:00:20:37:65:ce:11
  20:00:00:20:37:65:ce:22
  20:00:00:20:37:65:ce:33
  20:00:00:20:37:65:ce:44
  10:00:00:60:69:04:11:24
  10:00:00:60:69:04:11:23
  21:00:00:e0:8b:04:70:3b
  10:00:00:60:69:04:11:33
  20:00:00:20:37:65:ce:55
  20:00:00:20:37:65:ce:66
  00:00:00:00:00:00:00:00
}
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [1] 0
done ...
Update the FLASH? (yes, y, no, n): [yes] y
*Successfully saved the MS ACL to the flash.
switch:admin>

```

To delete a member from the ACL

1. Connect to the switch and log in as admin.
2. Enter the **msConfigure** command.
The command becomes interactive.
3. At the select prompt, enter **3** to delete a member based on its port/node WWN.
4. At the prompt, enter the WWN of the member to be deleted from the ACL.
5. At the prompt, enter **1** to verify the WWN you entered was deleted from the ACL.
6. After verifying that the WWN was deleted correctly, enter **0** at the prompt to end the session.
7. At the “Update the FLASH?” prompt, enter **y**.
8. Press **Enter** to update the nonvolatile memory and end the session.

Example

```

switch:admin> msconfigure
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [1] 3
Port/Node WWN (in hex): [00:00:00:00:00:00:00:00] 20:00:00:20:37:65:ce:aa
*WWN is successfully deleted from the MS ACL.
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [2] 1
MS Access List consists of (13): {
  20:00:00:20:37:65:ce:aa
  20:00:00:20:37:65:ce:bb
  20:00:00:20:37:65:ce:ff
  20:00:00:20:37:65:ce:11
  20:00:00:20:37:65:ce:22
  20:00:00:20:37:65:ce:33
  10:00:00:60:69:04:11:24
  10:00:00:60:69:04:11:23
  21:00:00:e0:8b:04:70:3b
  10:00:00:60:69:04:11:33
  20:00:00:20:37:65:ce:55
  20:00:00:20:37:65:ce:66
}
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [1] 0
done ...
Update the FLASH? (yes, y, no, n): [yes] y
*Successfully saved the MS ACL to the flash.
switch:admin>

```

Configuring the Server Database

The management server database can be viewed or cleared.

To view the contents of the management server database

1. Connect to the switch and log in as admin.
2. Enter the **msPlatShow** command.

The contents of the management server platform database are displayed.

```
switch:admin> msplatshow
-----
Platform Name: [9] "first obj"
Platform Type: 5 : GATEWAY
Number of Associated M.A.: 1
[35] "http://java.sun.com/products/plugin"
Number of Associated Node Names: 1
Associated Node Names:
10:00:00:60:69:20:15:71
-----
Platform Name: [10] "second obj"
Platform Type: 7 : HOST_BUS_ADAPTER
Number of Associated M.A.: 1
Associated Management Addresses:
[30] "http://java.sun.com/products/1"
Number of Associated Node Names: 1
Associated Node Names:
10:00:00:60:69:20:15:75
```

To clear the management server database

1. Connect to the switch and log in as admin.
2. Enter the **msplClearDb** command.
3. Enter **y** to confirm the deletion.

The management server platform database is cleared.

Controlling Topology Discovery

The topology discovery feature can be displayed, enabled, and disabled; it is disabled by default.

To display topology discovery status

1. Connect to the switch and log in as admin.
2. Enter the **mstdReadConfig** command.

```
switch:admin> mstdreadconfig
*MS Topology Discovery is Enabled.
switch:admin>
```

To enable topology discovery

1. Connect to the switch and log in as admin.
2. Enter the **mstdEnable** command to enable the discovery feature locally.

3. Enter the **mstdEnable all** command to enable the discovery feature on the entire fabric.

```
switch:admin> mstdenable

Request to enable MS Topology Discovery Service in progress....
*MS Topology Discovery enabled locally.

switch:admin> mstdenable ALL

Request to enable MS Topology Discovery Service in progress....
*MS Topology Discovery enabled locally.
*MS Topology Discovery Enable Operation Complete!!
```

To disable topology discovery

1. Connect to the switch and log in as admin.
2. Enter the **mstdDisable** command to disable the discovery feature locally.
A warning displays that all NID entries might be cleared.
3. Enter **y** to disable the discovery feature.
4. Enter the **mstdDisable all** command to disable the discovery feature on the entire fabric.
5. Enter **y** to disable the discovery feature.



Note

Disabling management server topology discover might erase all NID entries.

```
switch:admin> mstddisable

This may erase all NID entries. Are you sure? (yes, y, no, n): [no] y

Request to disable MS Topology Discovery Service in progress....
*MS Topology Discovery disabled locally.

switch:admin> mstddisable all

This may erase all NID entries. Are you sure? (yes, y, no, n): [no] y

Request to disable MS Topology Discovery Service in progress....
*MS Topology Discovery disabled locally.
*MS Topology Discovery Disable Operation Complete!!
```


Working With Diagnostic Features

This chapter provides information on diagnostics and how to display system, port, and specific hardware information. It also describes how to set up system logging mapping (**syslogd**) and how to set up the offloading of error messages (**supportSave**).

The purpose of the diagnostic subsystem is to evaluate the integrity of the system hardware.

Diagnostics are invoked two ways:

- Automatically during the power-on self test (POST)
- Manually using Fabric OS CLI commands

The error messages generated during these test activities are sent to the serial console and system message logs, whose output formats may differ slightly.

Use the **diagHelp** command to receive a list of all available diagnostic commands.

Refer to the *Fabric OS Command Reference Manual* for a complete description of each command.

Viewing Power-On Self Test

By default, when you power on the system, the boot loader automatically performs power-on self tests and loads a Fabric OS kernel image.

The POST tests provide a quick indication of hardware readiness when hardware is powered up. These tests do not require user input to function. They typically operate within several minutes, and support minimal validation because of the restriction on test duration. Their purpose is to give a basic health check before a new switch joins a fabric.

These tests are divided into two groups: POST1 and POST2. POST1 validates the hardware interconnect of the device, and POST2 validates the ability of the device to pass data frames between the ports. The specific set of diagnostic and test commands run during POST depends on the switch model.

POST1 cannot be bypassed and runs from the boot loader. The factory default configuration is also set to run POST2, but you can configure your switch to bypass POST2, which runs after the kernel image has started but before general system services such as login are enabled.

Although each test performed during POST2 is configurable, you should only modify a POST2 test if directed by your switch provider's customer service representative.

You can use the **diagDisablePost** command to disable POST2, and you can reenable it using the **diagEnablePost** command. Refer to the *Fabric OS Command Reference Manual* for additional information about these commands.

The following example shows a typical boot sequence, including POST messages:

```
The system is coming up, please wait...

Read board ID of 0x80 from addr 0x23
Read extended model ID of 0x16 from addr 0x22
Matched board/model ID to platform index 4
PCI Bus scan at bus 0
:   :   :
:   :   :
Checking system RAM - press any key to stop test

Checking memory address: 00100000

System RAM test using Default POST RAM Test succeeded.

Press escape within 4 seconds to enter boot interface.
Booting "Fabric Operating System" image.

Linux/PPC load:
BootROM command line: quiet
Uncompressing Linux...done.
Now booting the kernel
Attempting to find a root file system on hda2...
modprobe: modprobe: Can't open dependencies file /lib/modules/2.4.19/modules.dep (No such file or directory)
INIT: version 2.78 booting
INIT: Entering runlevel: 3
eth0: Link status change: Link Up. 100 Mbps Full duplex Auto (autonegotiation complete).

INITCP: CPLD Vers: 0x95 Image ID: 0x19
uptime: 2008; sysc_qid: 0
Fabric OS (Paulsa45)
Paulsa45 console login: 2005/03/31-20:12:42, [TRCE-5000], 0,, INFO, ?, trace:, trace_buffer.c, line: 1170

2005/03/31-20:12:42, [LOG-5000], 0,, INFO, SW4100_P45, Previous message repeat 1 time(s), trace_ulib.c, line: 540
2005/03/31-20:12:43, [HAM-1004], 219,, INFO, SW4100_P45, Processor rebooted - Unknown
SNMP Research SNMP Agent Resident Module Version 15.3.1.4
Copyright 1989, 1990, 1991, 1992, 1993, 1994, 1995, 1996, 1997, 1998, 1999, 2000, 2001 SNMP Research, Inc.
sysctrl: all services Standby
FSSK 2: chassis0(0): state not synchronized
FSSK 2: Services starting a COLD recovery
2005/03/31-20:12:48, [FSS-5002], 0,, INFO, SW4100_P45, chassis0(0): state not synchronized, svc.c, line: 318
2005/03/31-20:12:48, [FSS-5002], 0,, INFO, SW4100_P45, Services starting a COLD recovery, mdev.c, line: 638
2005/03/31-20:12:49, [MFIC-1002], 220,, INFO, Paulsa45, Chassis FRU header not programmed for switch NID, using
defaults (applies only to FICON environments).
sysctrl: all services Active
2005/03/31-20:12:50, [DGD-5001], 0,, INFO, SW4100_P45, Slot 0 has started POST., main.c, line: 1189
POST1: Started running Thu Mar 31 20:12:51 GMT 2005
POST1: Test #1 - Running turboramtest
POST1: Test #2 - Running portregtest
POST1: Script PASSED with exit status of 0 Thu Mar 31 20:12:54 GMT 2005 took (0:0:3)
POST2: Started running Thu Mar 31 20:12:55 GMT 2005
POST2: Test #1 - Running portloopbacktest (SERDES)
POST2: Test #2 - Running minicycle (SERDES)
POST2: Running diagshow
POST2: Script PASSED with exit status of 0 Thu Mar 31 20:13:12 GMT 2005 took (0:0:17)
2005/03/31-20:13:13, [BL-1000], 221,, INFO, Paulsa45, Initializing Ports... Enabling switch...
2005/03/31-20:13:13, [BL-1001], 222,, INFO, Paulsa45, Port Initialization Completed
2005/03/31-20:13:13, [EM-5012], 0,, INFO, SW4100_P45, EM: sent dumpready to ME., em.c, line: 2152
2005/03/31-20:13:13, [DGD-5002], 0,, INFO, SW4100_P45, Slot 0 has passed the POST tests., main.c, line: 936
```

If you choose to bypass POST2, or after POST2 completes, various system services are started and the boot process displays additional console status and progress messages.

Viewing Switch Status

Use the **switchStatusShow** command to display the overall status of the switch, including its power supplies, fans, and temperature. If the status of any one of these components is either marginal or down, the overall status of the switch is also displayed as marginal or down. If all components have a healthy status, the switch displays a healthy status.

To modify the rules used to classify the health of each component use the **switchStatusPolicySet** command. To view the rules, use the **switchStatusPolicyShow** command.

To view the overall status of the switch

1. Connect to the switch and log in as admin.
2. Enter the **switchStatusShow** command:

```
switch:admin> switchstatusshow
Switch Health Report                                Report time: 03/21/2005 03:50:36 PM
Switch Name:    SW3900
IP address:     10.33.54.176
SwitchState:    MARGINAL
Duration:       863:23
Power supplies monitor  MARGINAL
Temperatures monitor   HEALTHY
Fans monitor           HEALTHY
WWN monitor            HEALTHY
Standby CP monitor     HEALTHY
Blades monitor         HEALTHY
Flash monitor          HEALTHY
Marginal ports monitor HEALTHY
Faulty ports monitor   HEALTHY
Missing SFPs monitor   HEALTHY
All ports are healthy
switch:admin>
```

For more information on how the overall switch status is determined, refer to the **switchStatusPolicySet** command in the *Fabric OS Command Reference Manual*.

To display switch information

1. Connect to the switch and log in as admin.
2. At the command line, enter the **switchShow** command. This command displays the following information for a switch:
 - switchname - Displays the switch name.
 - switchtype - Displays the switch model and firmware version numbers.
 - switchstate - Displays the switch state: Online, Offline, Testing, or Faulty.
 - switchrole - Displays the switch role: Principal, Subordinate, or Disabled.
 - switchdomain - Displays the switch Domain ID.
 - switchid - Displays the embedded port D_ID of the switch.
 - switchwwn - Displays the switch World Wide Name.
 - switchbeacon - Displays the switch beaconing state: either ON or OFF.

The **switchShow** command also displays the following information for ports on the specified switch:

- Module type - The SFP type if a SFP is present.
- Port speed - The speed of the Port (1G, 2G, 4G, N1, N2, N4, or AN). The speed can be fixed, negotiated, or auto negotiated.
- Port state - The port status.
- Comment - Displays information about the port. This section might be blank or display WWN for F_Port or E_Port, Trunking state, upstream or downstream status.

The details displayed for each switch differ on different switch models. For more information refer to the **switchShow** command in the *Fabric OS Command Reference Manual*.

To display the uptime for a switch

1. Connect to the switch and log in as admin.
2. At the command line, enter the **uptime** command:

```
switch:admin> uptime
4:43am up 1 day, 12:32, 1 user, load average: 1.29, 1.31, 1.27
switch:admin>
```

The **uptime** command displays the length of time the system has been in operation, the total cumulative amount of uptime since the system was first powered-on, the date and time of the last reboot (applies only to FOS v3.x and v2.6.x systems), the reason for the last reboot (applies only to FOS v3.x and v2.6.x systems), and the load average over the past one minute (1.29 in the preceding example), five minutes (1.31 in the example), and 15 minutes (1.27 in the example). The reason for the last switch reboot is also recorded in the system message log.

Viewing Port Information

Use the commands that follow to view information about ports.

To view the status of a port

1. Connect to the switch and log in as admin.
2. Enter the **portShow** command, specifying the number that corresponds to the port you are troubleshooting. In this example, the status of port two is shown:

```
switch:admin> portshow 2
portName:
portHealth: HEALTHY

Authentication: None
portDisableReason: None
portCFlags: 0x1
portFlags: 0x4903      PRESENT ACTIVE E_PORT G_PORT U_PORT LOGICAL_ONLINE LOGIN
LED
portType: 10.0
portState: 1      Online
portPhys: 6      In_Sync
portScn: 16      E_Port      Trunk port
port generation number: 351
portId: 290200
portIfId: 43020004
portWwn: 20:02:00:05:1e:34:01:be
portWwn of device(s) connected:
      None
Distance: normal
portSpeed: N4Gbps

Interrupts:      0      Link_failure: 50      Frjt:      0
Unknown:      0      Loss_of_sync: 55      Fbsy:      0
Lli:      524      Loss_of_sig: 54
Proc_rgrd:      0      Protocol_err: 0
Timed_out:      0      Invalid_word: 0
Rx_flushed:      0      Invalid_crc: 0
Tx_unavail:      0      Delim_err: 0
Free_buffer:      0      Address_err: 0
Overrun:      0      Lr_in: 100
Suspended:      0      Lr_out: 50
Parity_err:      0      Ols_in: 50
2_parity_err:      0      Ols_out: 52
CMI_bus_err:      0

switch:admin>
```

Refer to the *Fabric OS Command Reference Manual* for additional **portShow** command information, such as the syntax for slot or port numbering.

To display the port statistics

1. Connect to the switch and log in as admin.
2. At the command line, enter the **portStatsShow** command.

Port statistics include information such as number of frames received, number of frames sent, number of encoding errors received, and number of class 2 and class 3 frames received.

Refer to the *Fabric OS Command Reference Manual* for additional **portStatsShow** command information, such as the syntax for slot or port numbering.

Example: portStatsShow command output

```

switch:admin> portstatsshow 3/7
stat_wtx          0          4-byte words transmitted
stat_wrx          0          4-byte words received
stat_ftx          0          Frames transmitted
stat_frx          0          Frames received
stat_c2_frx       0          Class 2 frames received
stat_c3_frx       0          Class 3 frames received
stat_lc_rx        0          Link control frames received
stat_mc_rx        0          Multicast frames received
stat_mc_to        0          Multicast timeouts
stat_mc_tx        0          Multicast frames transmitted
tim_rdy_pri       0          Time R_RDY high priority
tim_txcrd_z       0          Time BB credit zero
er_enc_in         0          Encoding errors inside of frames
er_crc            0          Frames with CRC errors
er_trunc          0          Frames shorter than minimum
er_toolong        0          Frames longer than maximum
er_bad_eof        0          Frames with bad end-of-frame
er_enc_out        0          Encoding error outside of frames
er_bad_os         0          Invalid ordered set
er_c3_timeout     0          Class 3 frames discarded due to timeout
er_c3_dest_unreach 0          Class 3 frames discarded due to destination unreachable
er_other_discard  0          Other discards
er_crc_good_eof   0          Crc error with good eof
er_inv_arb        0          Invalid ARB
open              0          loop_open
transfer          0          loop_transfer
opened           0          FL_Port opened
starve_stop       0          tenancies stopped due to starvation
fl_tenancy        0          number of times FL has the tenancy
nl_tenancy        0          number of times NL has the tenancy
zero_tenancy      0          zero tenancy

switch:admin>

```

To display a summary of port errors for a switch

1. Connect to the switch and log in as admin.
2. At the command line, enter the **portErrShow** command. Refer to the *Fabric OS Command Reference Manual* for additional portErrShow command information

Example: portErrShow command output

```

switch:admin> porterrshow
      frames  enc  crc  too  too  bad  enc  disc  link  loss  loss  frjt  fbsy
      tx   rx   in  err  shrt long  eof  out   c3 fail sync sig
sig=====
0:   22   24   0   0   0   0   0  1.5m  0   7   3   0   0   0
1:   22   24   0   0   0   0   0  1.2m  0   7   3   0   0   0
2:    0    0   0   0   0   0   0    0   0   0   0   0   0   0
3:    0    0   0   0   0   0   0    0   0   0   0   0   0   0
4:  149m  99m   0   0   0   0   0  448   0   7   6   0   0   0
5:  149m  99m   0   0   0   0   0  395   0   7   6   0   0   0
6:  147m  99m   0   0   0   0   0  706   0   7   6   0   0   0
7:  150m  99m   0   0   0   0   0  160   0   7   5   0   0   0
8:    0    0   0   0   0   0   0    0   0   0   0   0   0   0
9:    0    0   0   0   0   0   0    0   0   0   0   0   0   0
10:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
11:   0    0   0   0   0   0   0    0   0   0   0   2   0   0
12:   0    0   0   0   0   0   0    0   0   0   0   2   0   0
13:   0    0   0   0   0   0   0    0   0   0   0   2   0   0
14:   0    0   0   0   0   0   0    0   0   0   0   2   0   0
15:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
32:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
33:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
34:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
35:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
36:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
37:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
38:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
39:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
40:   99m  146m  0   0   0   0   0  666   0   6  796   7   0   0
41:   99m  149m  0   0   0   0   0  15k   0   2  303   4   0   0
42:   99m  152m  0   0   0   0   0  665   0   2  221   5   0   0
43:   99m  147m  0   0   0   0   0  16k   0   2  144   4   0   0
44:    0    0   0   0   0   0   0    0   0   0   0   0   0   0
45:    0    0   0   0   0   0   0    0   0   0   0   0   0   0
46:    0    0   0   0   0   0   0    0   0   0   0   2   0   0
47:    0    0   0   0   0   0   0    0   0   0   0   0   0   0

```

The **portErrShow** command output provides one output line per port. Refer to [Table 9-1](#) for a description of the error types.

Table 9-1 Error Summary Description

Error Type	Description
frames tx	Frames transmitted
frames rx	Frames received
enc in	Encoding errors inside frames
crc err	Frames with CRC errors
too shrt	Frames shorter than minimum
too long	Frames longer than maximum
bad eof	Frames with bad end-of-frame delimiters
enc out	Encoding error outside of frames
disc c3	Class 3 frames discarded

Table 9-1 Error Summary Description (Continued)

Error Type	Description
link fail	Link failures (LF1 or LF2 states)
loss sync	Loss of synchronization
loss sig	Loss of signal
frjt	Frames rejected with F_RJT
fbsy	Frames busied with F_BSY

Viewing Equipment Status

You can display status for fans, power supply, and temperature.



Note

The number of fans, power supply units, and temperature sensors depends on the switch type. For detailed specifications on these components, refer to the switch hardware reference manual.

The specific output from the status commands varies depending on the switch type.

To display the status of the fans

1. Connect to the switch and log in as admin.
2. Enter the **fanShow** command:

```
switch:admin> fanshow
Fan 1 is OK          speed is 7010 RPM
Fan 2 is OK          speed is 7180 RPM
Fan 3 is OK          speed is 7068 RPM
Fan 4 is OK          speed is 7116 RPM
Fan 5 is OK          speed is 7155 RPM
Fan 6 is OK          speed is 7001 RPM
switch:admin>
```

The possible status values are:

OK	Fan is functioning correctly.
Absent	Fan is not present.
Below minimum	Fan is present but rotating too slowly or stopped.
Above minimum	Fan is rotating too quickly.
Unknown	Unknown fan unit installed.
FAULTY	Fan has exceeded hardware tolerance.

To display the status of a power supply

1. Connect to the switch and log in as admin.
2. Enter the **psShow** command:

```
switch:admin> psshow
Power Supply #1 is OK
0335,FF2Z0007161,60-0000739-02, B,,DCJ3002-01P, B,FF2Z0007161
Power Supply #2 is faulty
0335,FF2Z0007176,60-0000739-02, B,,DCJ3002-01P, B,FF2Z0007176
switch:admin>
```

The possible status values are:

OK	Power supply functioning correctly.
Absent	Power supply not present.
Unknown	Unknown power supply unit installed.
Predicting failure	Power supply is present but predicting failure.
FAULTY	Power supply is present but faulty (no power cable, power switch turned off, fuse blown, or other internal error).

To display temperature status

1. Connect to the switch and log in as admin.
2. At the command line, enter the **tempShow** command:

```
switch:admin> tempshow
```

Index	Status	Centigrade	Fahrenheit
1	OK	21	70
2	OK	22	72
3	OK	29	84
4	OK	24	75
5	OK	25	77

```
switch:admin>
```

Information displays for each temperature sensor in the switch.

The possible temperature status values are:

OK	Temperature is within acceptable range.
FAIL	Temperature is outside of acceptable range. Damage might occur.

Viewing the System Message Log

The system message log feature enables messages to be saved across power cycles and reboots.

SilkWorm 12000, 24000, and 48000 director models maintain an independent system message log for each of the two CP blades. For these models, you should configure syslogd to support chronological system message logs. For details, see [“Configuring for syslogd” on page 9-12](#).

For details on error messages, refer to the *Fabric OS System Error Message Reference Manual*.

To display the system message log, with no page breaks

1. Connect to the switch and log in as admin.
2. Enter the **errDump** command at the command line.

To display the system message log one at a time

1. Connect to the switch and log in as admin.
2. Enter the **errShow** command at the command line.

To clear the system message log

1. Connect to the switch and log in as admin.
2. Enter the **errClear** command at the command line.

All switch and chassis events are removed.

Viewing the Port Log

The Fabric OS maintains an internal log of all port activity. The port log stores entries for each port as a circular buffer. Each port has space to store 8000 log entries. When the log is full, the newest log entries overwrite the oldest log entries. Port logs are not persistent and are lost over power-cycles and reboots. If the port log is disabled, an error message displays.



Note

Port log functionality is completely separate from the system message log. Port logs are typically used to troubleshoot device connections.

To view the port log

1. Connect to the switch and log in as admin.
2. Enter the **portLogShow** command:

```
switch:admin> portlogshow 12
```

time	task	event	port	cmd	args

Thu Apr 14 12:07:09 2005					
12:07:09.350	PORT	Rx	0	40	02ffffffd,00ffffffd,0608ffff,14000000
12:07:09.350	PORT	Tx	0	0	c0ffffffd,00ffffffd,060807fc
12:07:10.812	PORT	Tx	0	40	02ffffffd,00ffffffd,07feffff,14000000
12:07:10.813	PORT	Rx	0	0	c0ffffffd,00ffffffd,07fe0627
12:07:19.492	PORT	Tx	4	40	02ffffffd,00ffffffd,0800ffff,14000000
12:07:19.492	PORT	Tx	22	40	02ffffffd,00ffffffd,0802ffff,14000000
12:07:19.493	PORT	Rx	4	0	c0ffffffd,00ffffffd,08009287
12:07:19.493	PORT	Tx	24	40	02ffffffd,00ffffffd,0804ffff,14000000
12:07:19.494	PORT	Tx	31	40	02ffffffd,00ffffffd,0806ffff,14000000
12:07:19.494	PORT	Rx	22	0	c0ffffffd,00ffffffd,0802928d
12:07:19.494	PORT	Rx	24	0	c0ffffffd,00ffffffd,080492a3
12:07:19.495	PORT	Rx	31	0	c0ffffffd,00ffffffd,080692a7

Use the commands summarized in [Table 9-2](#) to view and manage port logs.

Table 9-2 Commands for Port Log Management

Command	Description
portLogClear	Clear port logs for all or particular ports.
portLogDisable	Disable port logs for all or particular ports.
portLogDump	Display port logs for all or particular ports, without page breaks.
portLogEnable	Enable port logs for all or particular ports.
portLogShow	Display port logs for all or particular ports, with page breaks.
Note: Refer to the <i>Fabric OS Command Reference Manual</i> for additional information about these commands.	

The **portLogDump** command output (trace) is a powerful tool that is used to troubleshoot fabric issues. The **portLogDump** output provides detailed information about the actions and communications within a fabric. By understanding the processes that are taking place in the fabric, issues can be identified and located.

The **portLogDump** command displays the port log, showing a portion of the Fibre Channel payload and header (FC-PH). The header contains control and addressing information associated with the frame. The payload contains the information being transported by the frame and is determined by the higher-level service or FC_4 upper level protocol. There are many different payload formats based on the protocol.

Because a **portLogDump** output is long, a truncated example is presented:

```
switch:admin> portlogdump
task event port cmd args
-----
16:30:41.780 PORT Rx 9 40 02ffffffd,00ffffffd,0061ffff,14000000
16:30:41.780 PORT Tx 9 0 c0ffffffd,00ffffffd,0061030f
16:30:42.503 PORT Tx 9 40 02ffffffd,00ffffffd,0310ffff,14000000
16:30:42.505 PORT Rx 9 0 c0ffffffd,00ffffffd,03100062
16:31:00.464 PORT Rx 9 20 02fffc01,00fffc0,0063ffff,01000000
16:31:00.464 PORT Tx 9 0 c0fffc0,00fffc01,00630311
16:31:00.465 nsd ctin 9 fc 000104a0,0000007f
16:31:00.465 nsd ctout 9 fc 00038002,00000003,01fffc01
16:31:00.466 PORT Tx 9 356 03fffc0,00fffc01,00630311,01000000
16:31:00.474 PORT Rx 9 0 c0fffc01,00fffc0,00630311
16:31:01.844 PORT Tx 9 40 02ffffffd,00ffffffd,0312ffff,14000000
16:31:01.854 PORT Rx 9 0 c0ffffffd,00ffffffd,03120064
16:31:01.963 PORT Rx 9 40 02ffffffd,00ffffffd,0065ffff,14000000
16:31:01.963 PORT Tx 9 0 c0ffffffd,00ffffffd,00650313
16:31:14.726 INTR pstate 0 LF2
16:31:14.729 PORT scn 0 137 00000000,00000000,00000008
16:31:14.729 PORT scn 0 129 00000000,00000000,00000400
16:31:14.729 PORT scn 0 2 00010004,00000000,00000002
16:31:14.730 SPEE sn 0 ws 00000002,00000000,00000000
<output truncated>
```

Configuring for syslogd

The system logging daemon (syslogd) is a process on UNIX, Linux, and some Windows systems that reads and logs messages as specified by the system administrator.

Fabric OS can be configured to use a UNIX-style syslogd process to forward system events and error messages to log files on a remote host system.

The host system can be running UNIX, Linux, or any other operating system that supports the standard syslogd functionality.



Note

Fabric OS releases earlier than v4.4.0 did not support UNIX local7 facilities; they supported kern facilities.

Starting with Fabric OS v4.4.0, kern facilities are no longer supported; UNIX local7 facilities are supported (the default facility level is 7).

Configuring for syslogd involves configuring the host, enabling syslogd on the SilkWorm model, and, optionally, setting the facility level.

Configuring the Host

Fabric OS supports a subset of UNIX-style message severities that default to the UNIX local7 facility. To configure the host, edit the `/etc/syslog.conf` file to map Fabric OS message severities to UNIX severities, as shown in [Table 9-3](#).

Table 9-3 Fabric OS to UNIX Message Severities

Fabric OS Message Severity	UNIX Message Severity
Critical (1)	Emergency (0)
Error (2)	Error (3)
Warning (3)	Warning (4)
Info (4)	Info (6)

In this example, Fabric OS messages map to local7 facility level 7 in the `/etc/syslog.conf` file:

```
local7.emerg      /var/adm/swcritical
local7.alert      /var/adm/alert7
local7.crit       /var/adm/crit7
local7.err        /var/adm/swerror
local7.warning    /var/adm/swwarning
local7.notice     /var/adm/notice7
local7.info       /var/adm/swinfo
local7.debug      /var/adm/debug7
```

If you prefer to map Fabric OS severities to a different UNIX local7 facility level, see [“To set the facility level” on page 9-14](#).

Configuring the Switch

Configuring the switch involves specifying syslogd hosts and, optionally, setting the facility level. You can also remove a host from the list of syslogd hosts.

To specify syslogd hosts

1. Connect to the switch and log in as admin.
2. Enter the **syslogDipAdd** command and specify an IP address.
3. Verify that the IP address was entered correctly, using the **syslogDipShow** command.

You can specify up to six host IP addresses for storing syslog messages, as shown in this example:

```
switch:admin> syslogdipadd 10.1.2.1
switch:admin> syslogdipadd 10.1.2.2
switch:admin> syslogdipadd 10.1.2.3
switch:admin> syslogdipadd 10.1.2.4
switch:admin> syslogdipadd 10.1.2.5
switch:admin> syslogdipadd 10.1.2.6
switch:admin> syslogdipshow
syslog.IP.address.1 10.1.2.1
syslog.IP.address.2 10.1.2.2
syslog.IP.address.3 10.1.2.3
syslog.IP.address.4 10.1.2.4
syslog.IP.address.5 10.1.2.5
syslog.IP.address.6 10.1.2.6
```

To set the facility level

1. Connect to the switch and log in as admin.
2. Enter the following command:

```
syslogdfacility -l n
```

n is a number from 0 through 7, indicating a UNIX local7 facility. The default is 7.

It is necessary to set the facility level only if you specified a facility other than local7 in the host */etc/syslog.conf* file.

To remove a syslogd host from the list

1. Connect to the switch and log in as admin.
2. Enter the **syslogDipRemove** command:

```
switch:admin> syslogdipremove 10.1.2.1
```

3. Verify the IP address was deleted using the **syslogDipShow** command.

Viewing and Saving Diagnostic Information

Enter the **supportShow** command to dump important diagnostic and status information to the session screen, where you can review it or capture its data.

To save a set of files that customer support technicians can use to further diagnose the switch condition, enter the **supportSave** command. The command prompts for an FTP server, packages the following files, and sends them to the specified server:

- the output of the **supportShow** command
- the contents of any trace dump files on the switch
- system message logs (for SilkWorm directors, **supportSave** saves the system message logs from both of the CP blades)

Refer also to [“Setting Up Automatic Trace Dump Transfers,”](#) next.

Setting Up Automatic Trace Dump Transfers

You can set up a switch so that diagnostic information is transferred automatically to a remote server. Then, if a problem occurs you can provide your customer support representative with the most detailed information possible. To ensure the best service, you should set up for automatic transfer as part of standard switch configuration, before a problem occurs.

Setting up for automatic transfer of diagnostic files involves the following tasks:

- Specify a remote server to store the files.
- Enable the automatic transfer of trace dumps to the server. (Trace dumps overwrite each other by default; sending them to a server preserves information that would otherwise be lost.)
- You should also set up a periodic checking of the remote server so that you are alerted if the server becomes unavailable and you can correct the problem.

After the setup is complete, you can run the **supportSave -c** command to save diagnostic information to the server (without specifying server details).

The following procedures describe in detail the tasks for setting up automatic transfer. For details on the commands, refer to the *Fabric OS Command Reference Manual*.

To specify a remote server

1. Verify that the FTP service is running on the remote server.
2. Connect to the switch and log in as admin.
3. Enter the following command:

```
supportftp -s
```

The command is interactive; provide the required information.

4. Respond to the prompts as follows:

<i>Host Name</i>	Enter the name or IP address of the server where the file is to be stored; for example, 192.1.2.3.
<i>User name</i>	Enter the user name of your account on the server; for example, "JohnDoe".
<i>Password</i>	Enter your account password for the server.
<i>Remote directory</i>	Specify a path name for the remote directory. Absolute path names can be specified using forward slash (/). Relative path names create the directory in the user's home directory on UNIX servers, and in the directory where the FTP server is running on Windows servers.

To enable the automatic transfer of trace dumps

1. Connect to the switch and log in as admin.
2. Enter the following command:

```
traceftp -e
```

To set up periodic checking of the remote server

1. Connect to the switch and log in as admin.
2. Enter the following command:

```
supportftp -t interval
```

The *interval* is in hours. The minimum interval is 1 hour. Specify 0 hours to disable the checking feature.

To save a comprehensive set of diagnostic files to the server

1. Connect to the switch and log in as admin.
2. Enter the following command:

```
supportsave -c
```

Troubleshooting should begin at the center of the SAN — the fabric. Because switches are located between the hosts and storage devices and have visibility into both sides of the storage network, starting with them can help narrow the search path. After eliminating the possibility of a fault within the fabric, see if the problem is on the storage side or the host side, and continue a more detailed diagnosis from there. Using this approach can quickly pinpoint and isolate problems.

For example, if a host cannot detect a storage device, run a switch command (such as **switchShow**) to find out if the storage device is logically connected to the switch. If not, focus first on the switch directly connecting to storage. Use your vendor-supplied storage diagnostic tools to better understand why it is not visible to the switch. If the storage can be detected by the switch, and the host still cannot detect the storage device, then there is still a problem between the host and switch.

This chapter provides information on troubleshooting and the most common procedures used to diagnose and repair issues. It also includes specific troubleshooting scenarios as examples.

Most Common Problem Areas

Refer to [Table 10-1](#) for a list of the most common problem areas that arise within SANs and a list of tools that can be used to resolve them.

Table 10-1 Common Troubleshooting Problems and Tools

Problem Area	Investigate	Tools
Fabric	<ul style="list-style-type: none">• Missing devices• Marginal links (unstable connections)• Incorrect zoning configurations• Incorrect switch configurations	<ul style="list-style-type: none">• Switch LEDs• Switch commands (for example, switchShow or nsAllShow) for diagnostics• Web or GUI-based monitoring and management software tools
Storage Devices	<ul style="list-style-type: none">• Physical issues between switch and devices• Incorrect storage software configurations	<ul style="list-style-type: none">• Device LEDs• Storage diagnostic tools• Switch commands (for example, switchShow or nsAllShow) for diagnostics

Table 10-1 Common Troubleshooting Problems and Tools (Continued)

Problem Area	Investigate	Tools
Hosts	<ul style="list-style-type: none"> • Incorrect host bus adapter installation • Incorrect device driver installation • Incorrect device driver configuration 	<ul style="list-style-type: none"> • Host adaptor LEDs • Host operating system diagnostic tools • Device driver diagnostic tools • Switch commands (for example, switchShow or nsAllShow) for diagnostics
Storage Management Applications	Incorrect installation and configuration of the storage devices that the software references. For example, if using a volume-management application, check for: <ul style="list-style-type: none"> • Incorrect volume installation • Incorrect volume configuration 	Application-specific tools and resources

Gathering Information for Technical Support

If you are troubleshooting a production system, you need to gather data quickly. As soon as a problem is observed, perform the following tasks (if using a dual CP system, run the commands on both CPs):

1. Enter the **supportSave** command to save RASLOG, TRACE, and supportShow (active CP only) information for the local CP to a remote FTP location. On a dual CP system, only the local CP information is saved and supportShow information is not available on the active CP.

For details about the supportSave command, refer to [“Viewing and Saving Diagnostic Information” on page 9-14](#) and [“Setting Up Automatic Trace Dump Transfers” on page 9-15](#)

2. Enter the **pdShow** command to display data from a panic dump file. The panic dump file contains information that might be useful to determine the cause of the system panic.
3. Enter the **saveCore** command to save or remove core files created by daemons.

For more details about these commands, refer to the *Fabric OS Command Reference Manual*.

Troubleshooting Questions

Common steps and questions to ask yourself when troubleshooting a system problem are as follows:

1. What is current Fabric OS level?
2. What is switch hardware version?
3. Is the switch operational?

4. Impact assessment and urgency:
 - Is the switch down?
 - Is it a standalone switch?
 - How large is the fabric?
 - Is the fabric redundant?
5. Run the **supportSave** command. (Refer to [“Viewing and Saving Diagnostic Information” on page 9-14](#) and [“Setting Up Automatic Trace Dump Transfers” on page 9-15](#)).
6. Document the sequence of events by answering the following questions:
 - What happened just prior to the problem?
 - Is the problem reproducible?
 - If so, what are the steps to produce the problem?
 - What configuration was in place when the problem occurred?
7. Did a failover occur?
8. Was security enabled?
9. Was POST enabled?
10. Are serial port (console) logs available?
11. Which CP blade was active? (only applicable to the SilkWorm 12000 or 24000)
12. What and when were the last actions or changes made to the system environment?

Use the following steps to retrieve as much of the following informational items as possible prior to contacting the SAN technical support vendor.

1. Switch information:
 - Serial number (located on the chassis)
 - World Wide Name (obtain using **licenseIdShow** or **wwn** commands)
 - Fabric OS version (obtain using **version** command)
 - Switch Configuration settings
 - supportSave output
 - pdShow and save Core output
2. Host information:
 - OS version and patch level
 - HBA type
 - HBA firmware version
 - HBA driver version
 - Configuration settings

3. Storage information:
 - Disk/tape type
 - Disk/tape firmware level
 - Controller type
 - Controller firmware level
 - Configuration settings
 - Storage software (such as EMC Control Center, Veritas SPC, etc.)

Analyzing Connection Problems

If a host is unable to detect its target (for example, a storage or tape device), you should begin troubleshooting the problem in the middle of the data path. Determine if the problem is *above* or *below* the starting point, then continue to divide the suspected problem path in half until you can pinpoint the problem.

Use the following procedures to analyze the problem:

To check the logical connection

1. Enter the **switchShow** command.
2. Review the output and determine if the device is logically connected to the switch:
 - A device that *is* logically connected to the switch will be registered as an F_Port or L_Port.
 - A device that is *not* logically connected to the switch will be registered as something other than an Nx_Port.
3. If the missing device *is* logically connected, proceed to the next troubleshooting procedure ([“To check the Simple Name Server \(SNS\)” on page 10-7](#)).
4. If the missing device is *not* logically connected, check the device and everything on that side of the data path. Also refer to [“Correcting Link Failures” on page 10-19](#).

This includes all aspects of the host OS, the driver settings and binaries, the device Basic Input Output System (BIOS) settings, the SFP, the cable going from the switch to the device, the SFP on the switch side of that cable, and all switch settings related to the device. Refer to [“To check for a loop initialization failure” on page 10-20](#) as the next potential trouble spot.

To check for fibre channel connectivity problems

1. Enter the **fcPing** command, which:
 - Checks the zoning configuration for the two ports specified.
 - Generates an ELS (Extended Link Service frame) ECHO request to the source port specified and validates the response.
 - Generates an ELS ECHO request to the destination port specified and validates the response.

2. Regardless of the device's zoning, the **fcPing** command sends the ELS frame to the destination port. A device can take any one of the following actions:

- Send an ELS Accept to the ELS request.
- Send an ELS Reject to the ELS request.
- Ignore the ELS request.

There are some devices that do not support the ELS ECHO request. In these cases, the device will either not respond to the request or send an ELS reject. When a device does not respond to the ELS request, further debugging is required; however, do not assume that the device is not connected to the Fibre Channel.

Following is sample output from the **fcPing** command in which one device accepts the request and another device rejects the request:

```
switch:admin> fcping 10:00:00:00:c9:29:0e:c4 21:00:00:20:37:25:ad:05
Source:      10:00:00:00:c9:29:0e:c4
Destination: 21:00:00:20:37:25:ad:05
Zone Check:  Not Zoned

Pinging 10:00:00:00:c9:29:0e:c4 [0x20800] with 12 bytes of data:
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1162 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1013 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1442 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1052 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1012 usec
5 frames sent, 5 frames received, 0 frames rejected, 0 frames timeout
Round-trip min/avg/max = 1012/1136/1442 usec

Pinging 21:00:00:20:37:25:ad:05 [0x211e8] with 12 bytes of data:
Request rejected
Request rejected
Request rejected
Request rejected
Request rejected
5 frames sent, 0 frames received, 5 frames rejected, 0 frames timeout
Round-trip min/avg/max = 0/0/0 usec
switch:admin>
```

Following is sample output from the **fcPing** command in which one device accepts the request and another device does not respond to the request:

```
switch:admin> fcping 0x020800 22:00:00:04:cf:75:63:85
Source:      0x020800
Destination: 22:00:00:04:cf:75:63:85
Zone Check:  Zoned

Pinging 0x020800 with 12 bytes of data:
received reply from 0x020800: 12 bytes time:1159 usec
received reply from 0x020800: 12 bytes time:1006 usec
received reply from 0x020800: 12 bytes time:1008 usec
received reply from 0x020800: 12 bytes time:1038 usec
received reply from 0x020800: 12 bytes time:1010 usec
5 frames sent, 5 frames received, 0 frames rejected, 0 frames timeout
Round-trip min/avg/max = 1006/1044/1159 usec

Pinging 22:00:00:04:cf:75:63:85 [0x217d9] with 12 bytes of data:
Request timed out
Request timed out
Request timed out
Request timed out
Request timed out
5 frames sent, 0 frames received, 0 frames rejected, 5 frames timeout
Round-trip min/avg/max = 0/0/0 usec
switch:admin>
```

For details about the **fcPing** command, refer to the *Fabric OS Command Reference Manual*.

To check the Simple Name Server (SNS)

1. Enter the **nsShow** command on the switch to which the device is attached:

```
The Local Name Server has 9 entries {
```

Type	Pid	COS	PortName	NodeName	TTL(sec)
*N	021a00;	2,3;	20:00:00:e0:69:f0:07:c6;10:00:00:e0:69:f0:07:c6;	895	
	Fabric Port Name: 20:0a:00:60:69:10:8d:fd				
NL	051edc;	3;	21:00:00:20:37:d9:77:96;20:00:00:20:37:d9:77:96;	na	
	FC4s: FCP [SEAGATE ST318304FC 0005]				
	Fabric Port Name: 20:0e:00:60:69:10:9b:5b				
NL	051ee0;	3;	21:00:00:20:37:d9:73:0f;20:00:00:20:37:d9:73:0f;	na	
	FC4s: FCP [SEAGATE ST318304FC 0005]				
	Fabric Port Name: 20:0e:00:60:69:10:9b:5b				
NL	051eel;	3;	21:00:00:20:37:d9:76:b3;20:00:00:20:37:d9:76:b3;	na	
	FC4s: FCP [SEAGATE ST318304FC 0005]				
	Fabric Port Name: 20:0e:00:60:69:10:9b:5b				
NL	051ee2;	3;	21:00:00:20:37:d9:77:5a;20:00:00:20:37:d9:77:5a;	na	
	FC4s: FCP [SEAGATE ST318304FC 0005]				
	Fabric Port Name: 20:0e:00:60:69:10:9b:5b				
NL	051ee4;	3;	21:00:00:20:37:d9:74:d7;20:00:00:20:37:d9:74:d7;	na	
	FC4s: FCP [SEAGATE ST318304FC 0005]				
	Fabric Port Name: 20:0e:00:60:69:10:9b:5b				
NL	051ee8;	3;	21:00:00:20:37:d9:6f:eb;20:00:00:20:37:d9:6f:eb;	na	
	FC4s: FCP [SEAGATE ST318304FC 0005]				
	Fabric Port Name: 20:0e:00:60:69:10:9b:5b				
NL	051eef;	3;	21:00:00:20:37:d9:77:45;20:00:00:20:37:d9:77:45;	na	
	FC4s: FCP [SEAGATE ST318304FC 0005]				
	Fabric Port Name: 20:0e:00:60:69:10:9b:5b				
N	051f00;	2,3;	50:06:04:82:bc:01:9a:0c;50:06:04:82:bc:01:9a:0c;	na	
	FC4s: FCP [EMC SYMMETRIX 5267]				
	Fabric Port Name: 20:0f:00:60:69:10:9b:5b				

2. Look for the device in the SNS list, which lists the nodes connected to that switch, allowing you to determine if a particular node is accessible on the network.
 - If the device is *not* present in the SNS, the problem is between the storage device and the switch. There might be a timeout communication problem between edge devices and the Name Server, or there might be a login issue. First check the edge device documentation to determine if there is a timeout setting or parameter that can be reconfigured. You should also check the port log for NS registration information and FCP probing failures (using the **fcProbeShow** command). If these queries do not help solve the problem, contact the support organization for the product that appears to be inaccessible.
 - If the device *is* listed in the SNS, the problem is between the storage device and the host. There might be a zoning mismatch or a host/storage issue. Proceed to [“To check for zoning problems” on page 10-8](#).
3. Enter the **portLoginShow** command to check the port login status.

4. Enter the **fcpProbeShow** command to display the FCP probing information for the devices attached to the specified F_Port or FL_Port. This information includes the number of successful logins and SCSI INQUIRY commands sent over this port and a list of the attached devices.
5. Check the port log to determine whether or not the device sent the FLOGI frame to the switch, and the switch probed the device.

To check for zoning problems

1. Enter the **cfgActvShow** command to determine if zoning is enabled.
If zoning is enabled, it is possible that the problem is being caused by zoning enforcement (for example, two devices in different zones cannot see each other).
2. Confirm that the specific edge devices that need to communicate with each other are in the same zone.
 - If they are in the same zone, perform the following tasks:
 - Enter the **portCamShow** command on the host port to verify that the target is present.
 - Enter the **portCamShow** command on the target.
 - Enter the **nsZoneMember** command on the host and target to determine whether or not the Name Server is aware that these devices are zoned together
 - If they are not in the same zone and zoning is enabled, proceed to [step 3](#).
3. Resolve zoning conflicts by putting the devices into the same zoning configuration.
Refer to [“Correcting Zoning Setup Issues” on page 10-10](#) for additional information.

Restoring a Segmented Fabric

Fabric segmentation is generally caused by:

- Incompatible fabric parameters (refer to [“To reconcile fabric parameters individually” on page 10-9](#)).
- Incorrect PID setting (refer to [“Configuring the PID Format” on page A-1](#)).
- Incompatible zoning configuration (refer to [“To check for zoning problems” on page 10-8](#)).
- Domain ID conflict (refer to [“To reconcile a domain ID conflict” on page 10-9](#)).
- A switch in a secure fabric is not running Secure Fabric OS.
Refer to the *Secure Fabric OS Administrator’s Guide* for additional information.

There are a number of settings that control the overall behavior and operation of the fabric. Some of these values, such as the domain ID, are assigned automatically by the fabric and can differ from one switch to another in the fabric. Other parameters, such as the BB credit, can be changed for specific applications or operating environments, but must be the same among all switches to allow the formation of a fabric.

The following fabric parameters must be identical for a fabric to merge:

- R_A_TOV
- E_D_TOV
- Data field size
- Sequence level switching

- Disable device probing
- Suppress class F traffic
- Per-frame route priority
- Long distance fabric (not necessary on Bloom-based fabrics)
- BB credit
- PID format

To reconcile fabric parameters individually

1. Log in to one of the segmented switches as admin (switch A).
2. Enter the **configShow** command.
3. Log in to another switch (switch B) in the same fabric as admin.
4. Enter the **configShow** command.
5. Compare the two switch configurations line by line and look for differences. Do this by comparing the two telnet windows or by printing the **configShow** output. Also verify that the fabric parameter settings (see [page 10-8](#)) are the same for *both* switches.
6. Connect to the segmented switch after the discrepancy is identified.
7. Disable the switch by entering the **switchDisable** command.
8. Enter the **configure** command to edit the fabric parameters for the segmented switch.
Refer to the *Fabric OS Command Reference Manual* for more detailed information.
9. Enable the switch by entering the **switchEnable** command.

Alternatively, you can reconcile fabric parameters by entering the **configUpload** command for each switch.

To download a correct configuration

You can restore a segmented fabric by downloading a previously saved correct backup configuration to the switch. Downloading in this manner reconciles any discrepancy in the fabric parameters and allows the segmented switch to rejoin the main fabric. For details on uploading and downloading configurations, refer to “[Maintaining Configurations](#)” on [page 4-1](#).

To reconcile a domain ID conflict

If a domain ID conflict appears, the conflict is only reported at the point where the two fabrics are physically connected. However, there might be several conflicting domain IDs, which will appear as soon as the initial conflict is resolved.

Typically, the fabric automatically resolves domain conflicts during fabric merges or builds unless Insistent Domain ID (IDID) is configured. If IDID is enabled, switches that cannot be programmed with a unique domain ID are segmented out. Check each switch that has IDID configured and make sure their domain IDs are unique within the configuration.

Repeat this procedure until all domain ID conflicts are resolved:

1. Enter the **fabricShow** command on a switch from one of the fabrics.
2. In a separate telnet window, enter the **fabricShow** command on a switch from the second fabric.

3. Compare the **fabricShow** output from the two fabrics. Note the number of domain ID conflicts; there might be several duplicate domain IDs that will need to be changed. Determine which switches have domain overlap and change the domain IDs for each of those switches.
4. Chose the fabric on which to change the duplicate domain ID; connect to the conflicting switch in that fabric.
5. Enter the **switchDisable** command.
6. Enter the **switchEnable** command.

This will enable the joining switch to obtain a new domain ID as part of the process of coming online. The fabric principal switch will allocate the next available domain ID to the new switch during this process.

7. Repeat [step 4](#) through [step 6](#) if additional switches have conflicting domain IDs.

Correcting Zoning Setup Issues

The types of zone configuration discrepancies that can cause segmentation are listed in [Table 10-2](#).

Table 10-2 Types of Zone Discrepancies

Conflict Cause	Description
Configuration mismatch	Occurs when zoning is enabled in both fabrics and the zone configurations that are enabled are different in each fabric.
Type mismatch	Occurs when the name of a zone object in one fabric is also used for a different type of zone object in the other fabric. A zone object is any device in a zone.
Content mismatch	Occurs when the definition in one fabric is different from the definition of a zone object with the same name in the other fabric.

[Table 10-3](#) summarizes commands that are useful for debugging zoning issues.

Table 10-3 Commands for Debugging Zoning

Command	Function
aliCreate	Use to create a zone alias.
aliDelete	Use to delete a zone alias.
cfgCreate	Use to create a zone configuration.
cfgShow	Displays zoning configuration.
licenseShow	Displays current license keys and associated (licensed) products.
switchShow	Displays currently enabled configuration and any E_Port segmentations due to zone conflicts.

Table 10-3 Commands for Debugging Zoning (Continued)

Command	Function
zoneAdd	Use to add a member to an existing zone.
zoneCreate	Use to create a zone. Before a zone becomes active, the cfgSave and cfgEnable commands must be used.
zoneHelp	Displays help information for zone commands.
zoneShow	Displays zone information.

Refer to [“Administering Advanced Zoning” on page 13-1](#) for additional information about setting up zoning properly. Also refer to the *Fabric OS Command Reference Manual* for details about zoning commands.

To correct a fabric merge problem quickly

You can correct zone conflicts by using the **cfgClear** command to clear the zoning database.



Caution

This is a disruptive procedure.

1. Determine which switch(es) have the incorrect zoning configuration; then, log in to the switches as admin.
2. Enter the **switchDisable** command.
3. Enter the **cfgDisable** command.
4. Enter the **cfgClear** command.



Caution

This command clears the zoning database on the affected switches.

5. Enter the **switchEnable** command. This forces a zone merge and populates the switches with the desired zoning database. The two fabrics will be merged again.

To correct a merge conflict without disrupting the fabric, first verify fabric merge problem, then edit zone configuration members, and then reorder the zone member list.

To verify a fabric merge problem

1. Enter the **switchShow** command to validate that the segmentation is due to a zone issue.
2. Refer to [Table 10-2 on page 10-10](#) to view the different types of zone discrepancies.

To edit zone configuration members

1. Log in to one of the switches in a segmented fabric as admin.
2. Enter the **cfgShow** command.
3. Print the output from the **cfgShow** command.
4. Start another telnet session and connect to the next fabric as an administrator.
5. Run the **cfgShow** command.

6. Print the output from the **cfgShow** command.
7. Compare the two fabric zone configurations line by line and look for an incompatible configuration.
8. Connect to one of the fabrics.
9. Run zone configure edit commands to edit the fabric zone configuration for the segmented switch (refer to [Table 10-3 on page 10-10](#) for specific commands).

If the zoneset members between two switches are not listed in the same order in both configurations, the configurations are considered a mismatch; this results in the switches being segmented in the fabric.

For example:

`[cfg1 = z1; z2]` is different from `[cfg1 = z2; z1]`, even though the members of the configuration are the same.

One simple approach to making sure that the zoneset members are in the same order is to keep the members in alphabetical order.

To reorder the zone member list

1. Use the output from the **cfgShow** for both switches.
2. Compare the order that the zone members are listed. Members must be listed in the same order.
3. Rearrange zone members so that the configuration for both switches is the same. Arrange zone members in alphabetical order, if possible.

Recognizing MQ-WRITE Errors

An MQ error is a message queue error. Identify an MQ error message by looking for the two letters *M* and *Q* in the error message:

```
2004/08/24-10:04:42, [MQ-1004], 218,, ERROR, ras007, mqRead, queue = raslog-test-
string0123456-raslog, queue I
D = 1, type = 2
```

MQ errors can result in devices dropping from the SNS or can prevent a switch from joining the fabric. MQ errors are rare and difficult to troubleshoot; resolve them by working with the switch supplier. When MQ errors are encountered, execute the **supportSave** command to capture debug information about the switch; then, forward the **supportSave** data to the switch supplier for further investigation.

Correcting I²C Bus Errors

I²C bus errors indicate defective hardware; the specific item is listed in the error message. Refer to the *Fabric OS System Error Message Reference Manual* for information specific to the error that was received. Some CPT and Environmental Monitor (EM) messages contain I²C-related information.

If the I²C message does not indicate the specific hardware that might be failing, begin debugging the hardware, as this is the most likely cause. The next sections provide procedures for debugging the hardware.

To check fan components

1. Log in to the switch as user.
2. Enter the **fanShow** command.
3. Check the fan status and speed output.

If any of the fan speeds display abnormal RPMs, replace the fan FRU.

To check the switch temperature

1. Log in to the switch as user.
2. Enter the **tempShow** command.
3. Check the temperature output.

Look for indications of high or low temperatures.

To check the power supply

1. Log in to the switch as user.
2. Enter the **psShow** command.
3. Check the power supply status. Refer to the appropriate hardware reference manual for details regarding the power supply status.

If any of the power supplies show a status other than OK, consider replacing the power supply as soon as possible.

To check the temperature, fan, and power supply

1. Log in to the switch as user.
2. Enter the **sensorShow** command. Refer to the *Fabric OS Command Reference Manual* for details regarding the sensor numbers.
3. Check the temperature output.

Look for indications of high or low temperatures.

4. Check the fan speed output.

If any of the fan speeds display abnormal RPMs, replace the fan FRU.

5. Check the power supply status.

If any of the power supplies show a status other than OK, consider replacing the power supply as soon as possible.

Correcting Device Login Issues

To try to pinpoint problems with device logins, use this procedure:

1. Log in to the switch as admin.
2. Enter the **switchShow** command; then, check for correct logins:

```
switch:admin> switchshow
switchName:      sw094135
switchType:      26.1
switchState:     Online
switchMode:      Native
switchRole:      Principal
switchDomain:     126
switchId:        fffc7e
switchWwn:       10:00:00:05:1e:34:00:69
zoning:          ON (cfg_em)
switchBeacon:    OFF
Port   Media Speed State
=====
  0    id   N1   Online   E-Port  10:00:00:60:69:11:f9:fc "2800_116"
  1    id   1G   Online   E-Port  10:00:00:60:69:11:f9:fc "2800_116"
  2    id   N2   No_Light
  3    id   2G   No_Light
  4    id   N2   Online   E-Port  (Trunk port, master is Port  5)
  5    id   N2   Online   E-Port  10:00:00:05:1e:34:00:8b "Dazzl25"
(downstream)(Trunk master)
  6    id   N2   No_Light
  7    id   N2   No_Light
  8    id   N1   Online   L-Port  4 public, 1 private, 1 phantom
  9    id   N2   No_Light
 10    id   N2   Online   G-Port
 11    id   N2   Online   F-Port  10:00:00:01:c9:28:c7:01
 12    id   N1   Online   L-Port  4 public, 1 private, 1 phantom
 13    --   N2   No_Module
 14    id   N2   Online   E-Port  (Trunk port, master is Port 15)
 15    id   N2   Online   E-Port  10:00:00:60:69:90:03:17 "TERM_113"
(downstream)(Trunk master)
```

3. Enter the **portCfgShow** command to see how the port is configured:

```
sw094135:admin> portcfgshow
Ports of Slot 0  0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Speed          AN 1G AN 2G  AN AN AN AN  AN AN AN AN  AN AN AN AN
Trunk Port      ON ON .. ON  ON ON ON ON  ON ON ON ON  ON ON ON ON
Long Distance   .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
VC Link Init    .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Locked L_Port   .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Locked G_Port   .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Disabled E_Port .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
ISL R_RDY Mode  .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Persistent Disable.. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Locked Loop HD  .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..

                                where AN:AutoNegotiate, ..:OFF, ??:INVALID.
p                                LM:L0.5
```

4. Enter the **portErrShow** command; then, check for errors that can cause login problems:

```
sw094135:admin> porterrshow
```

	frames tx	frames rx	enc in	crc err	too shrt	too long	bad eof	enc out	disc c3	link fail	loss sync	loss sig	frjt	fbsy
0:	38	75	0	0	0	0	0	0	0	9	11	0	0	0
1:	110	73	0	0	0	0	0	0	0	9	11	0	0	0
2:	0	0	0	0	0	0	0	38	0	4	0	2	0	0
3:	0	0	0	0	0	0	0	0	0	4	1	2	0	0
4:	59m	102	0	0	0	0	0	0	0	4	0	0	0	0
5:	59m	103	0	0	0	0	0	0	0	3	0	0	0	0
6:	0	0	0	0	0	0	0	21	0	3	0	0	0	0
7:	0	0	0	0	0	0	0	58	0	3	0	0	0	0
8:	81	19k	0	0	0	0	0	3.0m	0	5	43	0	0	0
9:	0	0	0	0	0	0	0	29	0	3	0	0	0	0
10:	12m	68m	0	0	0	0	0	13	43m	8	1	1	0	0
11:	30m	33m	0	0	0	0	0	0	0	8	1	1	0	0
12:	89	25k	0	0	0	0	0	2.9m	0	7	43	0	0	0
13:	0	0	0	0	0	0	0	0	0	3	0	0	0	0
14:	29m	82m	0	0	0	0	0	0	1.2m	4	1	1	0	0
15:	29m	81m	0	0	0	0	0	0	1.1m	4	1	1	0	0

- A high number of errors relative to the frames transmitted and frames received can indicate a marginal link (refer to [“Correcting Marginal Links”](#) on page 10-22 for additional information).
 - A steadily increasing number of errors can indicate a problem. Track errors by sampling the port errors every five or ten minutes.
5. Enter the **portFlagsShow** command; then, check to see how a port has logged in and where a login failed (if a failure occurred):

```
sw094135:admin> portflagsshow
```

Port	SNMP	Physical	Flags
0:	Online	In_Sync	PRESENT ACTIVE E_PORT G_PORT U_PORT LOGICAL_ONLINE
1:	Online	In_Sync	PRESENT ACTIVE E_PORT G_PORT U_PORT LOGICAL_ONLINE
2:	Offline	No_Light	PRESENT U_PORT LED
3:	Offline	No_Light	PRESENT U_PORT LED
4:	Online	In_Sync	PRESENT ACTIVE E_PORT G_PORT U_PORT LOGICAL_ONLINE
5:	Online	In_Sync	PRESENT ACTIVE E_PORT G_PORT U_PORT LOGICAL_ONLINE
6:	Offline	No_Light	PRESENT U_PORT LED
7:	Offline	No_Light	PRESENT U_PORT LED
8:	Online	In_Sync	PRESENT ACTIVE F_PORT L_PORT U_PORT LOGICAL_ONLINE
9:	Offline	No_Light	PRESENT U_PORT LED
10:	Online	In_Sync	PRESENT ACTIVE G_PORT U_PORT LOGIN LED
11:	Online	In_Sync	PRESENT ACTIVE F_PORT G_PORT U_PORT LOGICAL_ONLINE
12:	Online	In_Sync	PRESENT ACTIVE F_PORT L_PORT U_PORT LOGICAL_ONLINE
13:	Offline	No_Module	PRESENT U_PORT LED
14:	Online	In_Sync	PRESENT ACTIVE E_PORT G_PORT U_PORT LOGICAL_ONLINE
15:	Online	In_Sync	PRESENT ACTIVE E_PORT G_PORT U_PORT LOGICAL_ONLINE

6. Enter the **portLogDumpPort** *portid* command where the port ID is the port number; then, view the device to switch communication.

```
sw094135:admin> portlogdumpport 10
time          task          event  port  cmd  args
-----
12:38:21.590  SPEE          sn      10    WS   00000000,00000000,00000000
12:38:21.591  SPEE          sn      10    WS   000000ee,00000000,00000000
12:38:21.611  SPEE          sn      10    WS   00000001,00000000,00000000
12:38:21.871  SPEE          sn      10    NC   00000002,00000000,00000001
12:38:21.872  LOOP          loopscn 10    LIP   8002
12:38:22.171  LOOP          loopscn 10    TMO   2
12:38:22.171  INTR          pstate  10    LF2
12:38:22.172  INTR          pstate  10    OL2
12:38:22.172  INTR          pstate  10    LR3
12:38:22.172  INTR          pstate  10    AC
12:38:22.172  PORT          scn      10    11    00000000,00000000,00000002
12:38:22.311  PORT          scn      10    1     00000000,00000000,00000001
12:38:22.311  PORT          debug    10           00000001,00654320,00000001,00000000
12:38:22.311  PORT          debug    10           00000001,00654320,00000002,00000000
12:38:22.311  PORT          debug    10           00000001,00654320,00000003,00000000
12:38:22.313  PORT          Tx        10    164   02ffffff,00ffffff,025effff,10000000
12:38:22.314  PORT          debug    10           00000001,00654320,00000003,00000000 *
7
12:38:28.312  PORT          Tx        10    164   02ffffff,00ffffff,028fffff,10000000
12:38:34.312  PORT          Tx        10    164   02ffffff,00ffffff,0293ffff,10000000
12:38:40.312  PORT          Tx        10    164   02ffffff,00ffffff,0299ffff,10000000
12:38:46.312  PORT          Tx        10    164   02ffffff,00ffffff,029bffff,10000000
12:38:52.312  PORT          Tx        10    164   02ffffff,00ffffff,029dffff,10000000
12:38:58.312  PORT          Tx        10    164   02ffffff,00ffffff,02acffff,10000000
12:39:04.322  INTR          pstate  10    LR1
12:39:04.323  INTR          pstate  10    LR3
12:39:04.323  INTR          pstate  10    AC
12:39:04.323  PORT          scn      10    11    00000000,00000000,00000002
sw094135:root>
```



Note

Refer to [“Viewing the Port Log” on page 9-10](#) for overview information about a **portLogDump**. Refer to the *Brocade PortlogDump Reference Guide* for detailed information about decoding a **portLogDump**.

Identifying Media-Related Issues

This section provides procedures that help pinpoint any media-related issues in the fabric. The tests listed in [Table 10-4](#) are a combination of *structural* and *functional* tests that can be used to provide an overview of the hardware components and help identify media-related issues.

- *Structural* tests perform basic testing of the switch circuit. If a structural test fails, replace the main board or port blade.

- *Functional* tests verify the intended operational behavior of the switch by virtue of running frames through ports or bypass circuitry.

Table 10-4 Component Test Descriptions

Test Name	Operands	Checks
crossporttest	<code>[-nframes count]</code> <code>[-lb_mode mode]</code> <code>[-spd_mode mode]</code> <code>[-gbic_mode mode]</code> <code>[-norestore mode]</code> <code>[-ports itemlist]</code>	Functional test of port external transmit and receive path. The crossport is set to loopback using an external cable by default. However, this command can be used to check internal components by setting the <i>lb</i> operand to 5.
fporttest	<code>[-nframes count]</code> <code>[-ports itemlist]</code> <code>[-seed payload_pattern]</code> <code>[-width pattern_width]</code> <code>[-size pattern_size]</code>	Tests component to/from and HBA. Used to test online F_Port devices, N_Port devices, SFPs, and GBICs. Not supported for SilkWorm 3016 internal ports 1 through 14.
loopporttest	<code>[-nframes count]</code> <code>[-ports itemlist]</code> <code>[-seed payload_pattern]</code> <code>[-width pattern_width]</code>	Only tests components attached to a switch that are on a FC-AL.
spinfab	<code>[nMillionFrames [, ePortBeg [, ePortEnd [, setFail]]]]</code>	Tests components to/from a neighbor switch, such as ISLs, SFPs, and GBICs between switches.

The following procedures are for checking switch-specific components.

To test a port's external transmit and receive path

1. Connect to the switch and log in as admin.
2. Connect the port you want to test to any other switch port with the cable you want to test.
3. Enter the **crossPortTest** command with the following operands (this is a partial list. Refer to the *Fabric OS Command Reference Manual* for additional command information):

`[-nframes count]` Specify the number of frames to send.

`[-lb_mode mode]` Select the loopback point for the test.

`[-spd_mode mode]` Select the speed mode for the test.

`[-ports itemlist]` Specify a list of user ports to test.

Example

```
switch:admin> crossporttest
Running Cross Port Test .... passed.
```

To test a switch's internal components

1. Connect to the switch and log in as admin.
2. Connect the port you want to test to any other switch port with the cable you want to test.

3. Enter the **crossporttest -lb_mode 5** command where 5 is the operand that causes the test to be run on the internal switch components (this is a partial list—refer to the *Fabric OS Command Reference Manual* for additional command information):

[-nframes count] Specify the number of frames to send.

[-lb_mode mode] Select the loopback point for the test.

[-spd_mode mode] Select the speed mode for the test.

[-ports itemlist] Specify a list of user ports to test.

To test components to and from the HBA

1. Connect to the switch and log in as admin.
2. Enter the **fPortTest** command (refer to the *Fabric OS Command Reference Manual* for information on the command options). This command is not supported on the SilkWorm 3016 internal ports 1 through 14.

Example

```
switchname:admin> fporttest 100,8,0xaa55,2, 512
Will use pattern: aa55 aa55 aa55 aa55 aa55 aa55 ...
Running fPortTest .....
port 8 test passed.
value = 0
```

This example executed the **fPortTest** command 100 times on port 8 with payload pattern 0xaa55, pattern width 2 (meaning word width) and a default payload size of 512 bytes.

Refer to [Table 10-5](#) for a list of additional tests that can be used to determine the switch components that are not functioning properly. Refer to the *Fabric OS Command Reference Manual* for additional command information.

Table 10-5 Switch Component Tests

Test	Function
portloopbacktest	Functional test of port N to N path.
portregtest	A read and write test of the ASIC SRAMs and registers.
spinsilk	Functional test of internal and external transmit and receive paths at full speed.
sramretentiontest	Verifies that the data written into the miscellaneous SRAMs in the ASIC are retained after a 10-second wait.
crossporttest	Verifies the functional components of the switch.

Table 10-5 Switch Component Tests (Continued)

Test	Function
turboramtest	Verifies that the on chip SRAM located in the 2 Gbit/sec ASIC is using the Turbo-Ram BIST circuitry. These same SRAMs are tested by portregtest and sramretentiontest using PCI operations, but for this test the BIST controller is able to perform the SRAM write and read operations at a much faster rate.
statstest	Verifies that the ASIC statistics counter logic.
Related Switch Test Option:	
itemlist	Restricts the items to be tested to a smaller set of parameter values that you pass to the switch.

Correcting Link Failures

A link failure occurs when a server or storage is connected to a switch, but the link between the server/storage and the switch does not come up. This prevents the server/storage from communicating through the switch.

If the **switchShow** command or LEDs indicate that the link has not come up properly, use one or more of the following procedures.

To determine if the negotiation was successfully completed

The port negotiates the link speed with the opposite side. The negotiation usually completes in 1-2 seconds; however, sometimes the speed negotiation fails.



Note

Skip this procedure if the port speed is set to a static speed through the **portCfgSpeed** command.

1. Enter the **portCfgShow** command to display the port speed settings of all the ports.
2. Enter the **switchShow** command to determine if the port has module light.
3. Determine whether or not the port at 1 Gig/sec completes by entering the **portCfgSpeed** command. Then change the port speed to 2 Gig/sec. This should correct the negotiation by setting to one speed.
4. Enter the **portLogShow** or **portLogDump** command.
5. Check the events area of the output. The first example is 1 Gbit/sec and the second example is 2 Gbit/sec:

```
14:38:51.976  SPEE sn <Port#>  NC  00000001,00000000,00000001
```

```
14:39:39.227  SPEE          sn      <Port#>  NC  00000002,00000000,00000001
```

- *sn* indicates a speed negotiation.
- *NC* indicates negotiation complete.
- *01* or *02* indicate the speed that has been negotiated.

If these fields do not appear, proceed to the [step 6](#).

6. Correct the negotiation by entering the **portCfgSpeed** [slotnumber/]portnumber, speed_level command if the fields in [step 5](#) do not appear.

To check for a loop initialization failure

1. Verify the port is an L_Port.
 - a. Enter the **switchShow** command.
 - b. Check the comment field of the output to verify that the switch port indicates an L_Port. If a loop device is connected to the switch, the switch port must be initialized as an L_Port.
 - c. Check to ensure that the port state is online; otherwise, check for link failures.
2. Verify the loop initialization *if* the port is not an L_port.
 - a. Enter the **portLogShow** or **portLogDump** command.
 - b. Check argument number four for the *LISA frame* (0x11050100).

```
termB:admin> portlogdumpport 4
time          task          event  port cmd  args
-----
11:40:02.078  PORT          Rx3     23   20   22000000,00000000,ffffffff,11050100
Received LISA frame
```

The *LISA frame* indicates that the loop initialization is complete.

3. Skip point-to-point initialization.

The switch changes to point-to-point initialization after the Loop Initialization Soft Assigned (LISA) phase of the loop initialization. This behavior sometimes causes trouble with old HBAs. If this is the case, then:

Skip point-to-point initialization by using the **portCfgLport** Command.

To check for a point-to-point initialization failure:

1. Enter the **switchShow** command to confirm that the port is active and has a module that is synchronized.

If a fabric device or another switch is connected to the switch, the switch port must be online.
2. Enter the **portLogShow** or **portLogDump** commands.
3. Verify that the event area for the port state entry is *pstate*. The command entry *AC* indicates that the port has completed point-to-point initialization.

```
termB:root> portlogdumpport 4
time          task          event  port cmd  args
-----
11:38:21.726  INTR          pstate  4    AC
```

4. Skip over the loop initialization phase.

After becoming an active port, the port becomes an F_Port or an E_Port depending on the device on the opposite side. If the opposite device is a fabric device, the port becomes an F_Port. If the opposite device is another switch, the port becomes an E_Port.

Some fabric devices have problems with loop initialization. If this is evident, enter the **portCfgPort** *port #*, **1** command.

To correct a port that has come up in the wrong mode

1. Enter the **switchShow** command.
2. Refer to the comment fields (refer to [Table 10-6](#)) and follow the suggested actions.

Table 10-6 SwitchShow Output and Suggested Action

Output	Suggested Action
Disabled	Check the output from the switchShow command to determine whether or not the switch is disabled. If the port is disabled (for example, due to persistent disable or security reasons), attempt to resolve the issue and then enter the portEnable command.
Bypassed	Check the output from the switchShow command to determine whether or not the port is testing.
Loopback	Check the output from the switchShow command to determine whether or not the port is testing.
E_Port	If the opposite side is not another switch, the link has come up in a wrong mode. Check the output from the portLogShow or PortLogDump commands and identify the link initialization stage where the initialization procedure went wrong.
F_Port	If the opposite side of the link is a fabric device, the link has come up in a wrong mode. Check the output from portLogShow or PortLogDump commands.
G_Port	The port has not come up as an E_Port or F_Port. Check the output from portLogShow or PortLogDump commands and identify the link initialization stage where the initialization procedure went wrong.
L_Port	If the opposite side is <i>not</i> a loop device, the link has come up in a wrong mode. Check the output from portLogShow or PortLogDump commands and identify the link initialization stage where the initialization procedure went wrong.

Correcting Marginal Links

A marginal link involves the connection between the switch and the edge device. Isolating the exact cause of a marginal link involves analyzing and testing many of the components that make up the link (including the switch port, switch SFP, cable, the edge device, and the edge device SFP).

To troubleshoot a marginal link:

1. Enter the **portErrShow** command.

Example

```
switch:admin> porterrshow
```

	frames tx	frames rx	enc in	crc err	too shrt	too long	bad eof	enc out	disc c3	link fail	loss sync	loss sig	frjt	fbusy
0:	22	24	0	0	0	0	0	1.5m	0	7	3	0	0	0
1:	22	24	0	0	0	0	0	1.2m	0	7	3	0	0	0
2:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
3:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
4:	149m	99m	0	0	0	0	0	448	0	7	6	0	0	0
5:	149m	99m	0	0	0	0	0	395	0	7	6	0	0	0
6:	147m	99m	0	0	0	0	0	706	0	7	6	0	0	0
7:	150m	99m	0	0	0	0	0	160	0	7	5	0	0	0
8:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
9:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
10:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
11:	0	0	0	0	0	0	0	0	0	0	0	2	0	0
12:	0	0	0	0	0	0	0	0	0	0	0	2	0	0
13:	0	0	0	0	0	0	0	0	0	0	0	2	0	0
14:	0	0	0	0	0	0	0	0	0	0	0	2	0	0
15:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
32:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
33:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
34:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
35:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
36:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
37:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
38:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
39:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
40:	99m	146m	0	0	0	0	0	666	0	6	796	7	0	0
41:	99m	149m	0	0	0	0	0	15k	0	2	303	4	0	0
42:	99m	152m	0	0	0	0	0	665	0	2	221	5	0	0
43:	99m	147m	0	0	0	0	0	16k	0	2	144	4	0	0
44:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
45:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
46:	0	0	0	0	0	0	0	0	0	0	0	2	0	0
47:	0	0	0	0	0	0	0	0	0	0	0	0	0	0

```
switch:admin>
```

2. Establish if there are a relatively high number of errors (such as CRC errors or ENC_OUT errors), or if there are a steadily increasing number of errors to confirm a marginal link.
3. If you suspect a marginal link, isolate the areas by moving the suspected marginal port cable to a different port on the switch.

If the problem stops or goes away, the switch port or the SFP is marginal (proceed to [step 4](#))

If the problem does *not* stop or go away, refer to [step 7](#).

4. Replace the SFP on the marginal port.

5. You will need an adapter to run the loopback test for the SFP. Otherwise, run the **portLoopBack** test on the marginal port using the loopback mode *lb=5*. Refer to the *Fabric OS Command Reference Manual* for additional information.

Table 10-7 Loopback Modes

Loopback Mode	Description
1	Port Loopback (loopback plugs)
2	External (SERDES) loopback
5	Internal (parallel) loopback (indicates no external equipment)
7	Backend bypass & port loopback
8	Backend bypass & SERDES loopback
9	Backend bypass & internal loopback

6. Check the results of the loopback test and proceed as follows:
 - If the loopback test failed, the port is bad. Replace the port blade.
 - If the loopback test did not fail, the SFP was bad.
7. Optionally, to rule out cabling issues:
 - a. Insert a new cable in to the suspected marginal port.
 - b. Enter the **portErrShow** command to determine if a problem still exists.
 - If the **portErrShow** output displays a normal number of generated errors, the issue is solved.
 - If the **portErrShow** output still displays a high number of generated errors, follow the troubleshooting procedures for the Host or Storage device.

Inaccurate Information in the System Message Log

In rare instances, events gathered by the *track change* feature can report inaccurate information to the system message log.

For example, a user enters a correct user name and password, but the login was rejected because the maximum number of users had been reached. However, when looking at the system message log, the login was reported as successful.

If the maximum number of switch users has been reached, the switch will still perform correctly in that it will reject the login of additional users (even if they enter correct user name and password information).

However, in this limited example, the Track Change feature will report this event inaccurately to the system message log; it will appear that the login was successful. This scenario only occurs when the maximum number of users has been reached; otherwise, the login information displayed in the system message log should reflect reality.

For information regarding enabling and disabling track changes (TC), refer to [“Tracking and Controlling Switch Changes”](#).

Recognizing the Port Initialization and FCP Auto Discovery Process

The steps in the port initialization process represent a protocol used to discover the type of connected device and establish the port type. The possible port types are as follows:

U_Port	Universal FC port. This port type is the base Fibre Channel port type and all unidentified, or uninitiated ports are listed as U_Ports.
FL_Port	Fabric Loop port. This port connects both public and private loop devices.
G_Port	Generic port. This port acts a transition port for non-loop fabric capable devices (E_Port / F_Port).
E_Port	Expansion port. This port type is assigned to ISL links.
F_Port	Fabric port. This port is assigned to fabric capable devices.

The Brocade FCP auto discovery process enables private storage devices that accept PRLI to communicate in a fabric.

If device probing is enabled, the embedded port PLOGIs and attempts a PRLI into the device to retrieve information to enter into the Name Server. This enables private devices that do not FLOGI but accept PRLI to be entered in the Name Server and receive full fabric citizenship. Private devices that accept PRLI represent a majority of storage targets. Private hosts require the QuickLoop feature, which is not available in Fabric OS v4.0.0 or later.

A fabric-capable device will implicitly register information with Name Server during a FLOGI. These devices will typically register information with the Name Server before querying for a device list. The embedded port will still PLOGI and attempt PRLI with these devices.

You can view the Name Server table in Web Tools by clicking **Name Server** in the fabric toolbar. Refer to the *Web Tools Administrator's Guide* for more information.

Licensed Features

This section describes licensed Brocade Fabric OS features and includes the following chapters:

- [Chapter 11, *Administering Extended Fabrics*](#)
- [Chapter 12, *Administering ISL Trunking*](#)
- [Chapter 13, *Administering Advanced Zoning*](#)
- [Chapter 14, *Administering Advanced Performance Monitoring*](#)

Additional Brocade licensed features are documented in the following books:

- *Web Tools Administrator's Guide*
- *Fabric Manager Administrator's Guide*
- *Fabric Watch Administrator's Guide*
- *Secure Fabric OS Administrator's Guide*

Administering Extended Fabrics

This chapter contains procedures for using the Brocade Extended Fabrics licensed feature, which extends the distance that interswitch links (ISLs) can reach. To use extended ISL modes, you must first install the Extended Fabrics license. For details on obtaining and installing licensed features, refer to [“Maintaining Licensed Features” on page 2-8](#).

About Extended Link Buffer Allocation

As the distance between switches and the link speed increase, additional *buffer-to-buffer credits* are required to maintain maximum performance. The number of credits reserved for a port depends on the switch model and on the extended ISL mode for which it is configured.

SilkWorm 3016, 3250, 3850, 3900, 12000, 24000, and 48000 (FC2-16 Port Blades)

For the SilkWorm 3016, 3250, 3850, 3900, 12000, 24000, and 48000 using FC2-16 port blades, each *port group* contains four ports and uses a common pool of credits. Because the number of credits available for use within each port group is limited, configuring ports for extended links on these models might cause other ports to become disabled if there are not enough buffer credits available; for example:

- If two 2-Gbit/second ports in a group are configured for L1 mode, each will be allocated enough buffer-to-buffer credits to cause the other two ports in the group to become disabled.
- A port connected to a device that is in loopback mode might become disabled for lack of buffers if another port in that group is set to L2 mode.

Refer to [Chapter 5, “Configuring SilkWorm 12000, 24000, and 48000 Directors”](#) for details about port blade nomenclature.

SilkWorm 4012, 4100, and 48000 (FC4-16 and FC4-32 Port Blades)

For the SilkWorm 4100 and 48000 director using FC4-16 and FC4-32 port blades, buffer credits are used by all ports on chip. *Buffer-limited* port technology allows all ports to remain operational, even when extended links are in use.

For the SilkWorm 4012, buffer credits are available to all ports on the chip.

A buffer-limited port can come online with fewer buffer credits allocated than its configuration specifies, allowing it to operate at a reduced bandwidth instead of being disabled for lack of buffers.

Buffer-limited operation is supported for the L0 and LD extended ISL modes only, and is persistent across reboots, switch disabling and enabling, and port disabling and enabling.

Fabric Considerations

Consider these items that affect the fabric when you configure extended ISLs:

- The extended link configuration mode, L2 can reach 100 km at a speed of 2Gbit/sec between Brocade Fabric OS v4.x switches. However, it only supports a distance of up to 60 km if the link is established between Brocade Fabric OS v3.x and 4.x switches.
- The standard distance and long-distance ISL modes cannot be enabled at the same time.
- Balance the number of long-distance ISL connections and core-to-edge ISL connections within a switch. Configuring long-distance ISLs between core and edge switches is possible, but is not a recommended practice.
- Starting with Fabric OS v4.4.0, VC translation link initialization (an option of the **portCfgLongDistance** command) is enabled by default for long-distance links. For previous Fabric OS versions that support this option, it was disabled by default. To avoid inconsistency in the fabric, make sure that this value is enabled on both ends of the link. To connect to switches running Fabric OS versions earlier than v4.0.2 and v3.0.2c, make sure that VC translation link initialization is disabled because these versions do not support it.

Choosing an Extended ISL Mode

Table 11-1 lists the extended ISL modes for switches that have a Bloom ASIC. You can configure extended ISL modes with the **portCfgLongDistance** command when the Extended Fabrics license is activated.

Table 11-1 Extended ISL Modes: Switches with Bloom ASIC

Mode	Description	Buffer Allocation		Distance @ 1 Gbit/sec	Distance @ 2 Gbit/sec	Earliest Fabric OS Release	Extended Fabrics License Required?
		1 Gbit/sec	2 Gbit/sec				
L0	Level 0 static mode, the default.	5 (26) ^a	5 (26)	10 km	5 km	All	No
LE	Level E static mode, supports links beyond 5 km.	13	19	n/a	10 km	v3.x, v4.x	No
L0.5	Level 0.5 static mode (designated LM when listed with the portcfgshow command).	19	34	25 km	25 km	v3.1.0, v4.1.0, v4.x, 5.x	Yes
L1	Level 1 static mode.	27	54	50 km	50 km	All	Yes
L2	Level 2 static mode.	60	64	100 km	60 km	All	Yes
LD ^b	Dynamic mode uses automatic distance detection for a user-specified distance.	Auto	Auto	Auto	Auto	v3.1.0, v4.1.0, v4.4.0, 5.x (depending on the model)	Yes

a. For each data channel (in this case, there are 4) there are 5 credits, plus 6 extra credits.

b. The dynamic long-distance mode (LD) automatically configures the number of buffer credits required, based on the actual link distance.

Table 11-2 lists the extended ISL modes for switches that have a Goldeneye ASIC.

Table 11-2 Extended ISL Modes: Switches with Goldeneye ASIC

Mode	Buffer Allocation			Distance @ 1 Gbit/sec	Distance @ 2 Gbit/sec	Distance @ 4 Gbit/sec	Earliest Fabric OS Release	Extended Fabrics License Required?
	1 Gbit/sec	2 Gbit/sec	4 Gbit/sec					
L0	5 (26) ^a	5 (26)	2.5 km	10 km	5 km	2 km	All	No
LE	11	16	10 km	n/a	10 km	5 km	v3.x, v4.x	No

a. For each data channel (in this case, there are 4) there are 5 credits, plus 6 extra credits

Table 11-3 lists the extended ISL modes for switches that have a Condor ASIC.

Table 11-3 Extended ISL Modes: Switches with Condor ASIC

Mode	Buffer Allocation			Distance @ 1 Gbit/ sec	Distance @ 2 Gbit/ sec	Distance @ 4 Gbit/ sec	Earliest Fabric OS Release	Extended Fabrics License Required?
	1 Gbit/ sec	2 Gbit/ sec	4 Gbit/ sec					
L0	5 (26) ^a	5 (26)	2.5 km	10 km	5 km	2 km	All	No
LE	11	16	10 km	n/a	5 km	5 km	v3.x, v4.x	No
L0.5	18	31	56	25 km	25 km	25 km	v3.1.0, v4.1.0, v4.x, 5.x	Yes
L1	31	56	106	50 km	50 km	50 km	All	Yes
L12	56	106	206	100 km	100 km	100 km	All	Yes
LD ^b	Auto	Auto	Auto	Auto	Auto	Auto	v3.1.0, v4.1.0, v4.4.0, 5.x (depending on the model)	Yes

a. For each data channel (in this case, there are 4) there are 5 credits, plus 6 extra credits.

b. The dynamic long-distance mode (LD) automatically configures the number of buffer credits required, based on the actual link distance.

For dynamic long distance links, you can approximate the number of buffer credits using the following formula:

$$\text{Buffer credits} = [(\text{distance in km}) * (\text{data rate}) * 1000] / 2112$$

The data rate is 1.0625 for 1 Gbit/sec, 2.125 for 2 Gbit/sec, and 4.25 for 4 Gbit/sec and Fibre Channel. This formula provides the minimum number of credits that will be allocated to a given port; the actual number will likely be higher.

Configuring External Ports

The number of ports that can be configured per port group for each switch depends on both port speed and distance. The following sections specify the number of external ports available for each switch at a specific port speed and distance.

SilkWorm 3016

Extended links are supported on the two external ports (0 and 15). Configuring extended links on these ports might affect internal ports 9 and 10 because of shared buffers.

When configuring external ports for long distance, both port speed (1 or 2 Gbit/sec) and distance setting (L0.5, L1, L2, and LD) must be considered.

For external ports operating at 2 Gbit/sec, the following restrictions should be observed:

- Setting two ports to L2 (60 km) is not allowed.
- Setting one external port to L2 (60 km) and the other external port to L0.5 (25 km) disables two internal ports.

- If you set one external port to L2 (60 km) and the other external port is an E_Port, then one internal port is disabled.
- Setting both external ports to L1 (50 km) disables two internal ports.
- Setting one external port to L1 (50 km) and the other external port to L0.5 (25 km) disables one internal port.
- If you set one external port to L1 (50 km) and the other external port is an E_Port, then one internal port is disabled.
- Setting both external ports to L0.5 (25 km) disables one internal port.

SilkWorm 3250, 3850, 3900, 12000, 24000, and 48000 (FC2-16 port blades)

The number of ports that can be configured per port group at various distances is summarized in [Table 11-4](#).

Table 11-4 SilkWorm 3250, 3850, 3900, 12000, 24000, and 48000 (FC2-16 port blades)

Speed (Gbit/sec)	Number of Ports Allowed at Distance (km)					
	10 km	25 km	50 km	100 km	250 km	500 km
1	4 ports	4 ports	L1 mode: 4 ports LD mode: 3 ports	up to 2 ports	n/a	n/a
2	4 ports	3 ports	up to 2 ports	1 port	n/a	n/a
4	n/a	n/a	n/a	n/a	n/a	n/a

SilkWorm 4100

The number of ports that can be configured at various distances is summarized in [Table 11-5](#).

Table 11-5 SilkWorm 4100

Speed (Gbit/sec)	Number of Ports Allowed at Distance (km)					
	10 km	25 km	50 km	100 km	250 km	500 km
1	32 ports	32 ports	32 ports	up to 15 ports	up to 6 ports	up to 3 ports
2	32 ports	32 ports	up to 15 ports	up to 7 ports	up to 3 ports	n/a
4	32 ports	up to 15 ports	up to 7 ports	up to 3 ports	n/a	n/a

SilkWorm 4012

The number of ports that can be configured at various distances is summarized in [Table 11-6](#). Note that for the SilkWorm 4012, exact distances (rather than set, incremental distances) are used.

Table 11-6 SilkWorm 4012

Speed (Gbit/sec)	Number of Ports Allowed at Distance (km)			
	1 Port	Up to 2 Ports	Up to 3 Ports	Up to 4 ports
1	318 km	170 km	120.6 km	96 km
2	159 km	85 km	60.3 km	48 km
4	79.5 km	42.5 km	30.15 km	24 km

SilkWorm 48000 (FC4-16 port blades)

The number of ports that can be configured at various distances is summarized in [Table 11-7](#).

Table 11-7 SilkWorm 48000 (FC4-16 blades)

Speed (Gbit/sec)	Number of Ports Allowed at Distance (km)					
	10 km	25 km	50 km	100 km	250 km	500 km
1	16 ports	16 ports	16 ports	up to 12 ports	up to 4 ports	up to 2 ports
2	16 ports	16 ports	up to 12 ports	up to 5 ports	up to 2 ports	0
4	16 ports	up to 12 ports	up to 5 ports	up to 2 ports	0	0

SilkWorm 48000 (FC4-32 port blades)

The number of ports that can be configured at various distances is summarized in [Table 11-8](#).

Table 11-8 SilkWorm 48000 (FC4-32 blades)

Speed (Gbit/sec)	Number of Ports Allowed at Distance (km)					
	10 km	25 km	50 km	100 km	250 km	500 km
1	32 ports	32 ports	32 ports	up to 26 ports	up to 10 ports	up to 4 ports
2	32 ports	32 ports	up to 26 ports	up to 12 ports	up to 4 ports	0
4	32 ports	up to 26 ports	up to 12 ports	up to 6 ports	0	0

Configuring an Extended ISL

Before configuring the ISL, ensure that the following conditions are met:

- Extended ISL support for SilkWorm 2000-series switches is limited as follows:
 - Extended ISLs are not supported between SilkWorm 2000-series switches and other SilkWorm models.
 - To support extended ISLs between SilkWorm 2000-series switches, the **fabric.ops.mode.longDistance** parameter must be set to 1 on all switches in the fabric. Each switch must be disabled before setting this parameter.
- For fabrics that contain a mix of SilkWorm models, the **fabric.ops.mode.longDistance** parameter must be set to 0 (the default). Under certain circumstances (for example, if you want extended distance between Bloom-based switches) this mode needs to be enabled (set to 1) on switches running Fabric OS v3.x or v4.x. Talk to your switch provider for details.
- The ports on both ends of the ISL must have the same configuration.
- Use only qualified SFPs.

To configure an extended ISL

1. Connect to the switch and log in as admin.
2. If the fabric contains a mix of switches, use the **configure** command to make sure the fabric-wide configuration parameter **fabric.ops.mode.longDistance** is set to 0 on all switches in the fabric.

If the fabric contains SilkWorm 2000-series extended ISLs, use the **switchDisable** command to disable the switch and then use the **configure** command to set the fabric-wide configuration parameter **fabric.ops.mode.longDistance** to 1 on all switches in the fabric.

3. Enter the **portCfgLongDistance** command, using the following syntax:

```
portcfglongdistance [slotnumber/]portnumber [distance_level]
[vc_translation_link_init] [desired_distance]
```

<i>slotnumber</i>	Specify the slot number for SilkWorm 12000, 24000, and 48000 directors. This option is not applicable to fixed-port switches. The slot number must be followed by a slash (/) and the port number.
<i>portnumber</i>	Specify the port number.
<i>distance_level</i>	Specify the ISL mode to be set on the port (refer to Table 11-1 on page 11-3).

- vc_translation_link_init* This extended link initialization sequence, which is an enhanced link reset protocol, avoids excessive resetting of ports.
- By default this option is set to 1 (enabled).
- To prevent fabric segmentation, this option must be set to 0 (disabled) when connecting to switches running Fabric OS versions earlier than v3.0.2c or v4.0.2.
- It must be set to 1 (enabled) when configuring a trunk over extended fabrics.
- desired_distance* Required for a port configured for LD mode. Specify the desired distance, in kilometers, for the link. The specified value is the upper limit for calculating buffer availability for the port. If the measured distance is more than the specified *desired_distance*, the port is allocated the number of buffers required by the specified desired distance. (Fabric OS versions earlier than v4.4.0 do not support this parameter.)
- Note:** For SilkWorm 48000 with FC4-16 or FC4-32 port blades, ports will be shown as buffer limited.

4. Repeat [step 3](#) for the remote extended ISL port. Both the local and remote extended ISL ports must be configured to the same distance level. When the connection is initiated, the fabric will reconfigure.

The following example configures slot 1, port 1 for the LD link distance mode, enables the extended link initialization sequence, and sets the desired distance to 50 kilometers:

```
switch:admin> portcfglongdistance 1/1 LD 1 50
switch:admin>
```

Trunking Over Distance

Refer to [“Trunking Over Extended Fabrics” on page 12-9](#).

Administering ISL Trunking

This chapter contains procedures for using the Brocade ISL Trunking licensed feature, which optimizes the use of bandwidth by allowing a group of interswitch links to merge into a single logical link.

Overview

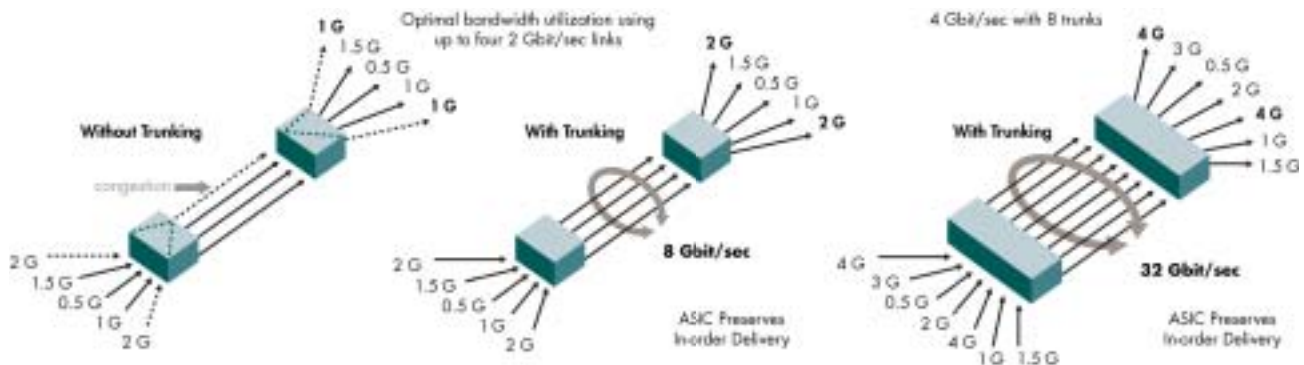
Brocade ISL Trunking reduces or eliminates situations that require static traffic routes and individual ISL management to achieve optimal performance. Trunking optimizes fabric performance by distributing traffic across the shared bandwidth of all the interswitch links in a trunking group, allowing traffic to flow through any available link in a group rather than restricting it to a specific, potentially congested link. The use of trunking results in simplified fabric design and management, lowered cost of ownership, and increased data availability.

To use trunking, you must first install the Brocade ISL Trunking license. For details on obtaining and installing licensed features, refer to [“Maintaining Licensed Features” on page 2-8](#). Trunking is enabled automatically when the ISL Trunking license is activated and ports are reinitialized (after installing the license, you enter the **switchDisable** and **switchEnable** commands), and trunks are easily managed using either Fabric OS CLI commands or Web Tools. You can enable and disable trunking and set trunk port speeds (for example, 2 Gig/sec, 4 Gig/sec, or autonegotiate) for entire switches or for individual ports.

Trunks distribute traffic dynamically and in order at the frame level, achieving greater performance with fewer interswitch links.

Trunks are compatible with both short wavelength (SWL) and long wavelength (LWL) fiber optic cables and transceivers.

[Figure 12-1](#) illustrates how trunking can result in more throughput by distributing data over four ISLs with no congestion. In a fabric that does not have trunking capability, some paths would be congested and other paths underutilized.

Figure 12-1 Distribution of Traffic over ISL Trunking Groups

Trunks operate best when the cable length of each trunked link is roughly equal to the others in the trunk. Cable lengths for participating links should differ by no more than 550 meters. For optimal performance, no more than 30 meters difference is recommended.

Connections between SilkWorm 4100 and SilkWorm 48000 (using FC4-16 and FC4-32 port blades) models support these advanced features:

- Up to eight ports in one trunk group to create high performance 32-Gbit/sec ISL trunks between switches
- ISL Trunking over longer distances than other models
- Dynamic trunk master reassignment if a trunk master is disabled (on other platforms, all ports on a trunk must be disabled temporarily to reassign a master)
- 4 Gbit/sec trunk links

The maximum number of ports per trunk and trunks per switch depends on the SilkWorm model.

For detailed information about trunking commands, refer to online help or the *Fabric OS Command Reference Manual*.

Standard Trunking Criteria

Observe the following criteria for standard distance trunking:

- There must be a direct connection between participating switches.
- Trunk ports must reside in the same port group.
- Trunk ports must run at the same speed (either 2 Gbit/sec or 4 Gbit/sec).
- Trunk ports must be set to the same ISL mode (L0 is the default). For details on extended ISL modes, refer to [Table 11-1 on page 11-3](#).
- Trunk ports must be E_Ports.
- Cable lengths for participating links should differ by no more than 550 meters.
- The **switch.interopMode** parameter must be set to 0. Refer to [“Configuring Interoperability Mode” on page B-1](#) for information and procedures related to interoperability mode.
- The port ISL mode must be disabled (using the **portCfgIslMode** command).

- **SilkWorm 3016:** Trunks are supported on the two external ports (0 and 15).

Fabric Considerations

The ISL Trunking feature is provided with the Fabric OS and can be activated by entering a license key, available from the switch supplier. When the ISL Trunking license is activated (after you have entered the **switchDisable** and **switchEnable** commands), trunking is automatically implemented for any eligible ISLs.

A license must be activated on each switch that participates in trunking. For the SilkWorm 12000, a single license key enables the feature on both logical switches.

To use ISL Trunking in the fabric, the fabric must be designed to allow trunking groups to form. To identify the most useful trunking groups, evaluate the traffic patterns before designing or redesigning the fabric. This also applies to the SilkWorm 24000 configured with two domains, and the SilkWorm 48000, which does not support two domains.

ISL Trunking can be used to simplify storage area network design and improve performance. When designing the SAN, consider the following recommendations in addition to the standard guidelines for SAN design:

- Evaluate the traffic patterns within the fabric.
- Place trunking-capable switches adjacent to each other.

This maximizes the number of trunking groups that can form. If you are using a core/edge topology, place trunking-capable switches at the core of the fabric and any switches that are not trunking-capable at the edge of the fabric.

- Activate an ISL Trunking license on each switch that is to participate in a trunking group.
- Cable lengths for participating links should differ by no more than 550 meters.
- When connecting two switches with two or more ISLs, ensure that all trunking requirements are met to allow a trunking group to form.
- Determine the optimal number of trunking groups between each set of linked switches, depending on traffic patterns and port availability.

The goal is to avoid traffic congestion without unnecessarily using ports that could be used to attach other switches or devices. Consider these points:

- Each physical ISL uses two ports that could otherwise be used to attach node devices or other switches.
- Trunking groups can be used to resolve ISL oversubscription if the total capability of the trunking group is not exceeded.
- Consider how the addition of a new path will affect existing traffic patterns:
 - A trunking group has the same link cost as the master ISL of the group, regardless of the number of ISLs in the group. This allows slave ISLs to be added or removed without causing data to be rerouted, because the link cost remains constant.
 - The addition of a path that is shorter than existing paths causes traffic to be rerouted through that path.
 - The addition of a path that is longer than existing paths might not be useful because the traffic will choose the shorter paths first.
- Plan for future bandwidth addition to accommodate increased traffic.

For trunking groups over which traffic is likely to increase as business requirements grow, consider leaving one or two ports in the group available for future nondisruptive addition of bandwidth.

- Consider creating redundant trunking groups where additional ports are available or paths are particularly critical.

This helps to protect against oversubscription of trunking groups, multiple ISL failures in the same group, and the rare occurrence of an ASIC failure.

- To provide the highest level of reliability, deploy trunking groups in redundant fabrics to further ensure ISL failures do not disrupt business operations.

Initializing Trunking on Ports

After you unlock the ISL Trunking license, you must reinitialize the ports being used for ISLs so that they recognize that trunking is enabled. This procedure only needs to be performed one time.

To reinitialize the ports, you can either disable and then reenable the switch, or disable and then reenable the affected ports.

To disable and reenable the switch

1. Connect to the switch and log in as admin.
2. Enter the **switchDisable** command.
3. Enter the **switchEnable** command.

To disable and reenable ports

1. Connect to the switch and log in as admin.
2. Enter the **portDisable** command. The format is:

portDisable *[slot/]port*

Slot is the slot number (SilkWorm 12000 and 24000 directors only) and *port* is the port number of the port you want to disable.

3. Enter the **portEnable** command. The format is:

portEnable *[slot/]port*

Slot is the slot number (SilkWorm 12000, 24000, and 48000 directors only) and *port* is the port number of the port you want to enable.

Monitoring Traffic

To implement ISL Trunking effectively, you must monitor fabric traffic to identify congested paths or to identify frequently dropped links. While monitoring changes in traffic patterns, you can adjust the fabric design accordingly, such as by adding, removing, or reconfiguring ISLs and trunking groups in problem areas.

There are three methods of monitoring fabric traffic:

- *Brocade Advanced Performance Monitoring* monitors traffic flow and allows you to view the impact of different fabric configurations on performance. Refer to [“Administering Advanced Performance Monitoring” on page 14-1](#) for additional information.
- *Brocade Fabric Watch* allows you to monitor traffic flow through specified ports on the switch and send alerts when the traffic exceeds or drops below configurable thresholds. Refer to the *Fabric Watch Administrator’s Guide* for additional information.
- Use the **portPerfShow** command as described in the following procedure to record traffic volume for each port in your fabric over time.

To use the portperfshow command

1. Connect to the switch and log in as admin.
2. Enter the following command:

```
portperfshow [interval]
```

where *interval* is the number of seconds between each data-gathering sample (the default is one sample every second).

3. Record the traffic flow for each port participating in an ISL.
4. Repeat [step 1](#) through [step 3](#) for each switch in the fabric until all ISL traffic flow is captured.

In a large fabric, it might be necessary to only identify and capture the key ISLs. However, you might want to continue this process throughout the day (or an entire work cycle), to capture varying traffic patterns under different conditions.

The following example shows a switch without trunking, and indicates that ports 0 through 2 are underutilized and ports 4 and 5 are congested:

```
switch:admin> portperfshow
0      1      2      3      4      5      6      7      Total
-----
0      0      0      145m   204m   202m   0      168m   719
0      0      0      145m   206m   208m   0      186m   745
switch:admin>
```

The following example shows traffic flowing through a trunking group (ports 5, 6, and 7). After port 6 fails, traffic is redistributed over the remaining two links in the group, ports 5 and 7:

```
switch:admin> portperfshow
0      1      2      3      4      5      6      7      Total
-----
0      0      0      0      0      145m   144m   145m   434
0      0      0      0      0      144m   143m   144m   431
0      0      0      0      0      162m   0      162m   324
0      0      0      0      0      186m   0      186m   372
0      0      0      0      0      193m   0      192m   385
0      0      0      0      0      202m   0      202m   404
0      0      0      0      0      209m   0      209m   418
switch:admin>
```

Enabling and Disabling ISL Trunking

You can enable or disable Brocade ISL Trunking for a single port or for an entire switch. When you execute the commands **portCfgTrunkPort** or **switchCfgTrunk** to update the trunking configuration, the ports for which the configuration applies are disabled and reenabled with the new trunk configuration. As a result, traffic through those ports could be disrupted.

To enable or disable ISL Trunking on one port

1. Connect to the switch and log in as admin.
2. Enter the **portCfgTrunkPort** command. The format is:

```
portcfgtrunkport [slotnumber/]portnumber mode
```

<i>slotnumber</i>	Specifies the number of the slot in which the port blade containing the port is located (this operand only required for switches with slots such as the SilkWorm 12000, 24000, and 48000 directors).
<i>portnumber</i>	Specifies the number of the port on which you want to enable or disable trunking.
<i>mode</i>	Enables (1) or disables (0) trunking on the specified port.

The following example enables trunking on slot 1, port 3:

```
switch:admin> portcfgtrunkport 1/3 1
done.
switch:admin>
```

To enable or disable ISL Trunking for all of the ports on a switch

1. Connect to the switch and log in as admin.
2. Enter the **switchCfgTrunk** command. The format is:

```
switchcfgtrunk mode
```

Mode 1 enables and mode 0 disables ISL Trunking for all ports on the switch.

The following example enables trunking all ports in the switch.

```
switch:admin> switchcfgtrunk 1
Committing configuration...done.
switch:admin>
```

Setting Port Speeds

For long-distance ports, if a port is set to autonegotiate port speed, the maximum speed (which is 4 Gbit/sec) is assumed for reserving buffers for the port—this wastes buffers if the port is actually running at 2 Gbit/sec. For long-distance ports, it is best to set the port speed (this applies to SilkWorm 4100 and SilkWorm 48000 only).

You can set the port speed for one port or for an entire switch. Trunked ports must be set to the same speed.

To set the speed for one port

1. Connect to the switch and log in as admin.
2. Enter the **portCfgSpeed** command. The format is:

```
portcfgspeed [slotnumber/]portnumber, speed_level
```

- | | |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>slotnumber</i> | For bladed systems only, specify the slot number of the port to be configured, followed by a slash (/). This operand is only required for switches with slots, such as the SilkWorm 12000, 24000, and 48000. |
| <i>portnumber</i> | Specifies the port number relative to its slot for bladed systems. |
| <i>speedlevel</i> | Specifies the speed of the link: <ul style="list-style-type: none"> • 0—Autonegotiating mode. The port automatically configures for the highest speed. • 1—one Gbit/sec mode. Fixes the port at a speed of one Gbit/sec. Changing the speed to one Gbit/sec causes the port to be excluded from the trunk group. • 2—two Gbit/sec mode. Fixes the port at a speed of two Gbit/sec. • 4—four Gbit/sec mode. Fixes the port at a speed of four Gbit/sec. (SilkWorm 200E, 4100 and 48000 only.) |

The following example sets the speed for port 3 on slot 2 to two Gbit/sec:

```
switch:admin> portcfgspeed 2/3 2
done.
switch:admin>
```

The following example sets the speed for port 3 on slot 2 to autonegotiate:

```
switch:admin> portcfgspeed s/3 0
done.
switch:admin>
```

To set the speed for all of the ports on the switch

1. Connect to the switch and log in as admin.
2. Enter the **switchCfgSpeed** command. The format is:

```
switchcfgspeed speedlevel
```

speedlevel Specifies the speed of the link:

- 0—Auto-negotiating mode. The port automatically configures for the highest speed.
- 1—one Gbit/sec mode. Fixes the port at a speed of one Gbit/sec. Changing the speed to one Gbit/sec causes the port to be excluded from the trunk group.
- 2—two Gbit/second mode. Fixes the port at a speed of two Gbit/second.
- 4—four Gbit/second mode. Fixes the port at a speed of four Gbit/second. (SilkWorm 200E, 4100 and 48000 switches only.)

The following example sets the speed for all ports on the switch to two Gbit/second:

```
switch:admin> switchcfgspeed 2
Committing configuration...done.
switch:admin>
```

The following example sets the speed for all ports on the switch to autonegotiate:

```
switch:admin> switchcfgspeed 0
Committing configuration...done.
switch:admin>
```

Displaying Trunking Information

The **trunkShow** command offers an efficient means of listing out all the trunks and members of a trunk. You can easily discover the peer ports for disabling a port, disconnecting or adding additional members. Viewing trunks is also useful for when you want to ensure that trunks are formed correctly.

Use the **trunkShow** command to display the following information about ISL Trunking groups:

- Number identifier.
- Port-to-port connections, listed in the following format: *local port number -> remote port number*.
- WWNs of the remote switches.
- Deskew values (the time difference, in nanoseconds divided by 10, for traffic to travel over each ISL as compared to the shortest ISL in the group). The system automatically sets the minimum deskew value of the shortest ISL to 15.
- Master ports.

To display trunking information:

1. Connect to the switch and log in as admin.
2. Enter the **trunkShow** command.

This example shows three trunking groups (1, 2, and 3); ports 1, 4, and 14 are masters:

```
switch:admin> trunkshow
1: 1 -> 1    10:00:00:60:69:04:10:83    deskew 16 Master
   0 -> 0    10:00:00:60:69:04:10:83    deskew 15
2: 4 -> 4    10:00:00:60:69:04:01:94    deskew 16 Master
   5 -> 5    10:00:00:60:69:04:01:94    deskew 15
   7 -> 7    10:00:00:60:69:04:01:94    deskew 17
   6 -> 6    10:00:00:60:69:04:01:94    deskew 16
3: 14 -> 14   10:00:00:60:69:04:10:83    deskew 16 Master
   15 -> 15   10:00:00:60:69:04:10:83    deskew 15
switch:admin>
```

Trunking Over Extended Fabrics

In addition to the criteria listed in [“Standard Trunking Criteria” on page 12-2](#), observe the following criteria for trunking over extended fabrics:

- ISL Trunking over extended fabrics is supported on switches running Fabric OS v3.2.0 (or later) or v4.4.0 (or later).
- Extended Fabrics and ISL Trunking licenses are required on all participating switches.
- The **vc_translation_link_init** parameter must be set the same on all ports in an extended trunk. (For details on this parameter, see [page 11-8](#).)

Troubleshooting Trunking Problems

If you have difficulty with trunking, try the solutions in this section.

Listing Link Characteristics

If a link that is part of an ISL Trunk fails, use the **trunkDebug** command to troubleshoot the problem, as shown in the following procedure:

1. Connect to the switch and log in as admin.
2. Enter the following command:

```
trunkDebug port, port
```

port Specifies the number of a port in an ISL Trunking group.

The **trunkDebug** command displays the possible reason that two ports cannot be trunked, including the following reasons:

- The switch does not support trunking.
- A trunking license is required.

- Trunking is not supported in switch interoperability mode.
- Port trunking is disabled.
- The port is not an E_Port.
- The port is not 2 Gbit/sec or 4 Gbit/sec.
- The port connects to different switches.
- The ports are not same speed, or they are not set to a valid speed.
- The ports are not set to the same long distance mode.
- Local or remote ports are not in same port group.
- The difference in the cable length among trunked links is greater than the allowed difference.

This example shows that port 3 is not configured as an E_Port:

```
switch:admin> trunkdebug 3 5
port 3 is not E port
switch:admin>
```

Recognizing Buffer Underallocation

For SilkWorm 3250, 3850, 3900, 12000, and 24000 models, if there is an underallocation or overcommitment of buffers to ports configured for extended trunking, the switches at both ends of the trunk try to disable some ports so that others can operate using the available buffers. (Standard trunks are not affected by buffer allocation.)

This issue of buffer underallocation does not apply to the SilkWorm 4100 and SilkWorm 48000 models.

A port disabled at one end because of buffer underallocation causes all the disabled ports at the other end to become enabled. Some of these enabled ports become disabled due to a lack of buffers, which in turn triggers ports to be enabled once again at the other end. While the system is stabilizing the buffer allocation, it warns that ports are disabled due to lack of buffers, but it does not send a message to the console when buffers are enabled. The system requires a few passes to stabilize the buffer allocation. Ultimately, the number of ports for which buffers are available come up and stabilize. You should wait for stabilization, and then proceed with correcting the buffer allocation situation.

To get out of buffer-limited mode on E_ports or LD_ports:

1. Change the LD/L1/L2/L0.5 port speed to a lower speed (of non-buffer limited ports).
2. Change the LD port's estimated distance to a shorter distance (of non-buffer limited ports).
3. Change LD/L1/L2/L0.5 back to L0 (of non-buffer limited ports).
4. If you are in buffer-limited mode on the LD port, then increase the estimated distance.

These changes are implemented only after disabling (portDisable) and enabling (portEnable) the buffer-limited port (or buffer-limited switch).

Reconfiguring a port to LD from another mode can result in the port being disabled for lack of buffers—this does not apply to the SilkWorm 4100 and SilkWorm 48000 (using FC4-16 and FC4-32 port blades). If this happens:

- In Fabric OS v4.2.x, reconfigure the disabled LD port back to the original mode.

- In Fabric OS v4.4.0 and later, specify a slightly shorter distance for the *desired_distance* parameter in the **portCfgLongDistance** command.

Administering Advanced Zoning

This chapter provides procedures for using the Brocade Advanced Zoning feature.

Zoning Terminology

The following terms are used in the Advanced Zoning procedures:

- **zone**
A *zone* is a region within the fabric where a specified group of fabric-connected devices (called *zone members*) have access to one another. When zoning is enabled, objects not explicitly defined in a zone are isolated, and members in the zoned fabric do not have access to them.
- **zone configuration**
A group of one or more zones is called a *zone configuration*.
- **defined zone configuration**
The complete set of all zone members defined in a fabric is called the *defined zone configuration*.
- **effective zone configuration**
Zoning procedures change zone objects in the defined configuration. When you enable a configuration with the **cfgEnable** command, it becomes the *effective zone configuration*. The effective zone configuration is restored after a switch reboot.
- **saved zone configuration**
A copy of the defined zone configuration (plus the name of the effective zone configuration) can be saved with the **cfgSave** command. The resulting *saved zone configuration* is restored after a switch reboot. If you make changes to the defined zone configuration but do not save them, there will be differences between the defined zone configuration and the saved zone configuration.

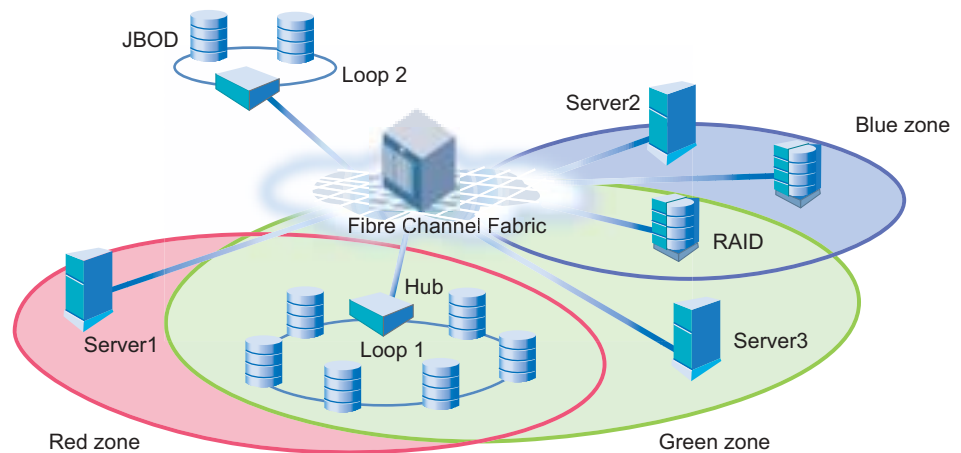
To use zoning, you must install Brocade Advanced Zoning licenses on all the switches in the fabric before attempting to bring a switch into the fabric. If a Zoning license is removed, you must make sure it is reinstalled properly on the affected switch before attempting the **cfgEnable** zoning operation. Failure to follow these steps can cause inconsistency of the zoning configuration on the affected switches should a zoning operation be attempted from a remote switch in the fabric. On the affected switches, an error message indicates that the Zoning license is missing.

You can use zones to logically consolidate equipment for efficiency or to facilitate time-sensitive functions; for example, to create a temporary zone to back up nonmember devices.

Any zone object connected to the fabric can be included in one or more zones. Zone objects can communicate only with other objects in the same zone. For example, consider [Figure 13-1](#), which shows:

- Three zones are configured, named Red, Green, and Blue.
- Server 1 can communicate only with the Loop 1 devices.
- Server 2 can communicate only with the RAID and Blue zone devices.
- Server 3 can communicate with the RAID device and the Loop1 device.
- The Loop 2 JBODs are not assigned to a zone; no other zoned fabric device can access them.

Figure 13-1 Zoning Example



To list the commands associated with zoning, use the **zoneHelp** command.

For detailed information on the zoning commands used in the procedures, refer to the *Fabric OS Command Reference Manual* or to the online man page for each command.

Zoning Concepts

Before using the procedures, you should become familiar with the zoning concepts described in the following sections.

Zone Types

[Table 13-1](#) summarizes the types of zoning.

Table 13-1 Types of Zoning

Zone Type	Description
Storage-based	Storage units typically implement LUN-based zoning, also called <i>LUN masking</i> . LUN-based zoning limits access to the LUNs on the storage port to the specific WWN of the server HBA. It is needed in most SANs. It functions during the probe portion of the SCSI initialization. The server probes the storage port for a list of available LUNs and their properties. The storage system compares the WWN of the requesting HBA to the defined zone list, and returns the LUNs assigned to the WWN. Other LUNs on the storage port are not made available to the server.
Host-based	Host-based zoning can implement WWN or LUN masking.
Fabric-based	<p>Fabric switches implement fabric-based zoning, in which the zone members are identified by WWN or port location in the fabric. Fabric-based zoning is also called <i>name server-based</i> or <i>soft zoning</i>.</p> <p>Brocade switches might also provide additional hardware enforcement of the zone. When a device queries the fabric name server, the name server determines the zones in which the device belongs. The server returns information on all members of the zones in the fabric to the device. Devices in the zone are identified by node WWN, port WWN, or domain, port of the switch to which the device is connected.</p> <p>Fabric-based zoning is perhaps the most controversial aspect of zoning. There are several approaches for implementing fabric zoning; all will work, in most cases. However, there are pros and cons to each form. The primary forms are summarized in Table 13-2.</p>

Table 13-2 Approaches to Fabric-Based Zoning

Zoning by:	
Single HBA	Zoning by single HBA most closely re-creates the original SCSI bus. Each zone created has only one HBA (initiator) in the zone; each of the target devices is added to the zone. Typically, a zone is created for the HBA and the disk storage ports are added. If the HBA also accesses tape devices, a second zone is created with the HBA and associated tape devices in it. In the case of clustered systems, it could be appropriate to have an HBA from each of the cluster members included in the zone; this is equivalent to having a shared SCSI bus between the cluster members and presumes that the clustering software can manage access to the shared devices. In a large fabric, zoning by single HBA requires the creation of possibly hundreds of zones; however, each zone contains only a few members. Zone changes affect the smallest possible number of devices, minimizing the impact of an incorrect zone change. This zoning philosophy is the preferred method.
Application	Zoning by application typically requires zoning multiple, perhaps incompatible, operating systems into the same zones. This method of zoning creates the possibility that a minor server in the application suite could disrupt a major server (such as a Web server) disrupting a data warehouse server. Zoning by application can also result in a zone with a large number of members, providing greater susceptibility to administrative errors, such as registered state change notifications (RSCNs) going out to a larger group than necessary.
Operating system	Zoning by operating system has issues similar to zoning by application. In a large site, this type of zone can become very large and complex. When zone changes are made, they typically involve applications rather than a particular server type. If members of different operating system clusters can see storage assigned to another cluster, they might attempt to own the other cluster's storage and compromise the stability of the clusters.
Port allocation	Avoid zoning by port allocation unless the administration team has very rigidly enforced processes for port and device allocation in the fabric. It does, however, provide some positive features. For instance, when a storage port, server HBA, or tape drive is replaced, the change of WWN for the new device is of no consequence. As long as the new device is connected to the original port, it continues to have the same access rights. The ports on the edge switches can be pre-associated to storage ports, and control of the fan-in ratio (the ratio of the input port to output port) can be established. With this pre-assigning technique, the administrative team cannot overload any one storage port by associating too many servers with it.
No fabric zoning	Using no fabric zoning is the least desirable zoning option because it allows devices to have unrestricted access on the fabric. Additionally, any device attached to the fabric, intentionally or maliciously, likewise has unrestricted access to the fabric. This form of zoning should only be utilized in a small and tightly controlled environment, such as when host-based zoning or LUN masking is deployed.

Zone Objects

A *zone object* is any device in a zone, such as the:

- Physical port number or area ID on the switch
- Node World Wide Name (N-WWN)
- Port World Wide Name (P-WWN)

Zone objects identified by port number or area number are specified as a pair of decimal numbers in the form *d, area* (*d* is the domain ID of the switch and *area* is the area number on that switch).

For example, on SilkWorm 12000 or 24000 models, “4, 46” specifies port 14 in slot number 3 (domain ID 4, area 46). On fixed-port models, “3,13” specifies port 13 in switch domain ID 3.

When the physical port number specifies a zone object, then all devices connected to that port are in the zone. If the physical port is an arbitrated loop, then all devices on the loop are part of the zone.

World Wide Names are specified as 8-byte (16-digit) hexadecimal numbers, separated by colons: for example, 10:00:00:90:69:00:00:8a. When a node name specifies a zone object, all ports on such a device are in the zone. When a port name specifies a zone object, only the single port is in the zone.

The types of zone objects used to define a zone can be mixed and matched. For example, a zone defined with the zone objects 2,12; 2,14; 10:00:00:80:33:3f:aa:11 contains the devices connected to domain 2, ports 12 and 14, and a device with the WWN (either node name or port name) 10:00:00:80:33:3f:aa:11 that is connected on the fabric.

Zone Aliases

A *zone alias* is a name assigned to a device or a group of devices. By creating an alias, you can assign a familiar name to a device or group multiple devices into a single name. This simplifies cumbersome data entry and allows an intuitive naming structure (such as using “NT_Hosts” to define all NT hosts in the fabric).

Zone aliases also simplify repetitive entry of zone objects such as port numbers or a WWN. For example, you can use the name “Eng” as an alias for “10:00:00:80:33:3f:aa:11”.

A useful convention is to name zones for the initiator they contain. For example, if you use the alias SRV_MAILSERVER_SLT5 to designate a mail server in PCI slot 5, then the alias for the associated zone is ZNE_MAILSERVER_SLT5. This clearly identifies the server host bus adapter (HBA) associated with the zone.

Zone configuration naming is more flexible. One configuration should be named PROD_*fabricname*, where *fabricname* is the name that the fabric has been designated. The purpose of the PROD configuration is to easily identify the configuration that can be implemented and provide the most generic services. If other configurations are used for specialized purposes, names such as “BACKUP_A,” “RECOVERY_2,” and “TEST_18jun02” can be used.

Zone Configurations

A *zone configuration* is a group of one or more zones. A zone can be included in more than one zone configuration. When a zone configuration is in effect, all zones that are members of that configuration are in effect.

The different types of zone configurations are:

- **Defined Configuration.** The complete set of all zone objects defined in the fabric.
- **Effective Configuration.** A single zone configuration that is currently in effect. The effective configuration is built when an administrator enables a specified zone configuration.
- **Saved Configuration.** A copy of the defined configuration plus the name of the effective configuration, which is saved in flash memory by the **cfgSave** command. (You can also use the **configUpload** command to provide a backup of the zoning configuration and the **configDownload** command to restore the zoning configuration.) There might be differences between the saved configuration and the defined configuration if the system administrator has modified any of the zone definitions and has not saved the configuration.
- **Disabled Configuration.** The effective configuration is removed from flash memory.

On power-up, the switch automatically reloads the saved configuration. If a configuration was active when it was saved, the same configuration is reinstated on the local switch with an autorun of the **cfgEnable** command.

You can establish a zone by identifying zone objects using one or more of the following *zoning schemes*:

- **Domain, port number level.** All members are specified by *domain ID*, *port number*, or *domain, area number* pair or aliases, described in [“Zone Aliases” on page 13-5](#).
- **World Wide Name (WWN) level.** All members are specified only by World Wide Name (WWNs) or aliases of WWNs. They can be node or port versions of the WWN.
- **Mixed zoning.** A zone containing members specified by a combination of *domain, port number*, and/or *domain, area number* and WWN.

Zoning Enforcement

Software-enforced and hardware-enforced zoning are supported.

Software-Enforced Zoning

Zoning enables users to restrict access to devices in a fabric. Software-enforced zoning prevents hosts from discovering unauthorized target devices, while hardware-enforced zoning prevents a host from accessing a device it is not authorized to access.

Software-enforced zoning:

- Is also called *soft zoning*, *Name Server zoning*, *fabric-based zoning*, *session-based zoning*, or *hardware-assisted zoning*.
- Is available on 1 Gbit/sec, 2 Gbit/sec, and 4 Gbit/sec platforms.
- Prevents hosts from discovering unauthorized target devices.
- Ensures that the name server does not return any information to an unauthorized initiator in response to a name server query.
- Is always active whenever a zone configuration is in effect.
- Does not prohibit access to the device. If an initiator has knowledge of the network address of a target device, it does not need to query the Name Server to access it, which could lead to undesired access to a target device by unauthorized hosts.

- Is exclusively enforced through selective information presented to end nodes through the fabric Simple Name Server (SNS). When an initiator queries the name server for accessible devices in the fabric, the name server returns only those devices that are in the same zone as the initiator. Devices that are not part of the zone are not returned as accessible devices.

Hardware-Enforced Zoning

Hardware-enforced zoning is specified without using the mixed zoning scheme (mixed zones contain domains, ports and WWNs as zone members). Brocade switches augment software-enforced zoning with hardware enforcement. The exact methodology varies on different switch models.

Hardware-enforced zoning (also called *hard zoning*):

- Prevents a host from accessing a device it is not authorized to access.
- Checks each frame before it is delivered to a zone member and discards it if there is a zone mismatch. When hardware-enforced zoning is active, the Brocade switch monitors the communications and blocks any frames that do not comply with the effective zone configuration. The switch performs this blocking at the transmit side of the port on which the destination device is located.
- Is enforced at the ASIC level. Each ASIC maintains a list of source port IDs that have permission to access any of the ports on that ASIC.

Fabric OS uses hardware-enforced zoning (on a per-zone basis) whenever the fabric membership or zone configuration changes.

Table 13-3 shows the various Brocade switch models, the hardware zoning methodology for each, and tips for best usage.

Table 13-3 Enforcing Hardware Zoning

Fabric Type	Methodology	Best Practice
SilkWorm 2000-series	<p>Enables hardware-enforced zoning only on domain, port zones; WWN or mixed zones are not hardware-enforced. Any domain, port zone that overlaps a mixed or WWN zone is not hardware-enforced.</p> <p>An overlap occurs when a member specified by WWN is connected to a port in a domain, port zone. The domain, port zone loses its hardware enforcement even though a review of the zone configuration does not indicate it.</p>	Use <i>domain, port</i> identifiers. Do not identify a zone member by its WWN.
SilkWorm 200E, SilkWorm 3000-series, SilkWorm 4100, SilkWorm 12000, 24000, and 48000 models	<p>Enable hardware-enforced zoning on domain, port zones, and WWN zones. Overlap of similar zone types does not result in the loss of hardware enforcement. Overlap with other zone type results in the loss of hardware enforcement.</p> <p>As in the SilkWorm 2000-series switches, connecting a device specified by WWN into a port specified in a domain, port zone results in loss of the hardware enforcement in both zones.</p>	Use either WWN or <i>domain, port</i> identifiers.
Mixed switches	<p>Enable hardware-enforced zoning according to each switch type. Use the portZoneShow command to find the zone type to which a device is attached.</p>	<p>Use <i>domain, port</i> identifiers.</p> <p>You can use WWN identifiers if you place disk and tape targets on SilkWorm 3000-series, and 12000 and 24000 models, and do not use <i>domain, port</i> identifiers.</p>

Figure 13-2 shows a fabric with four nonoverlapping hardware-enforced zones.

Figure 13-2 Hardware-Enforced Nonoverlapping Zones

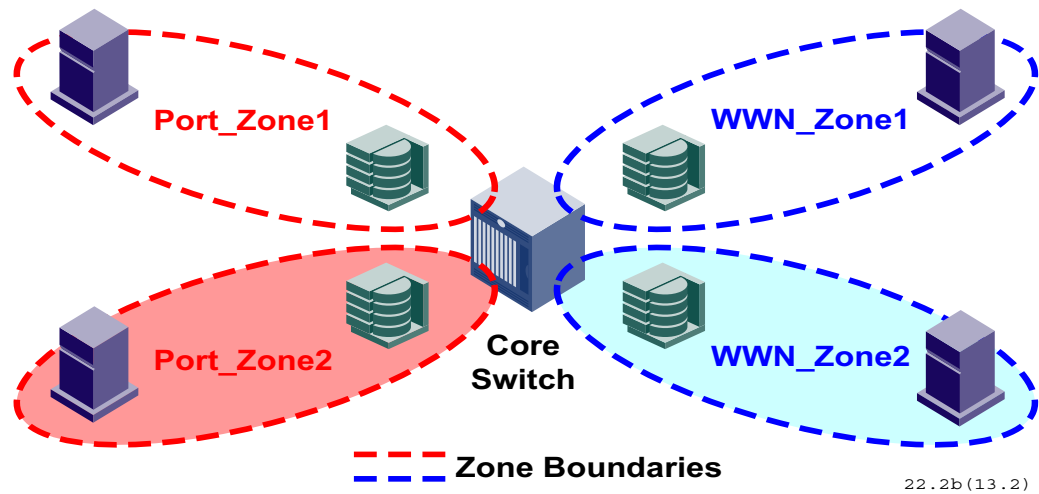
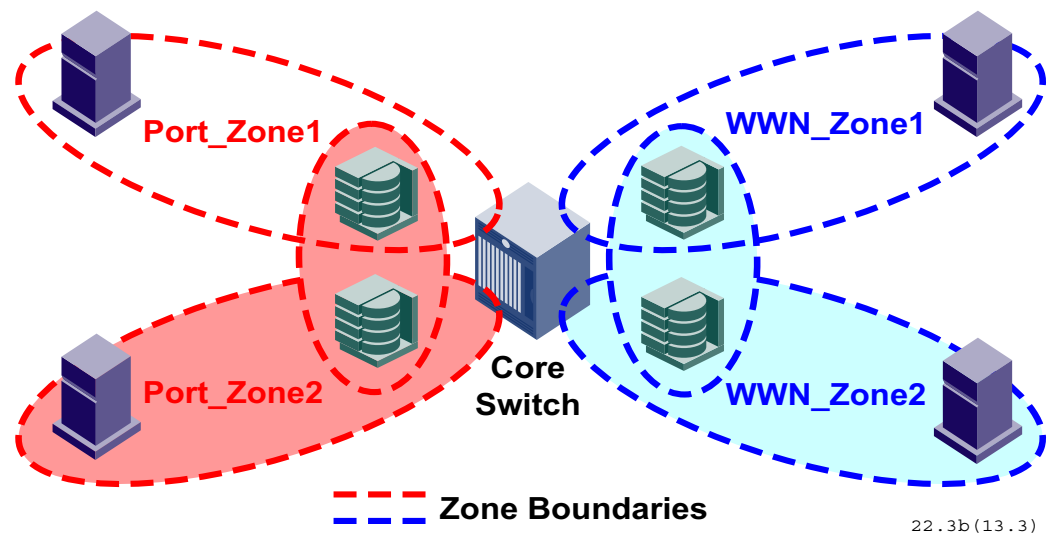
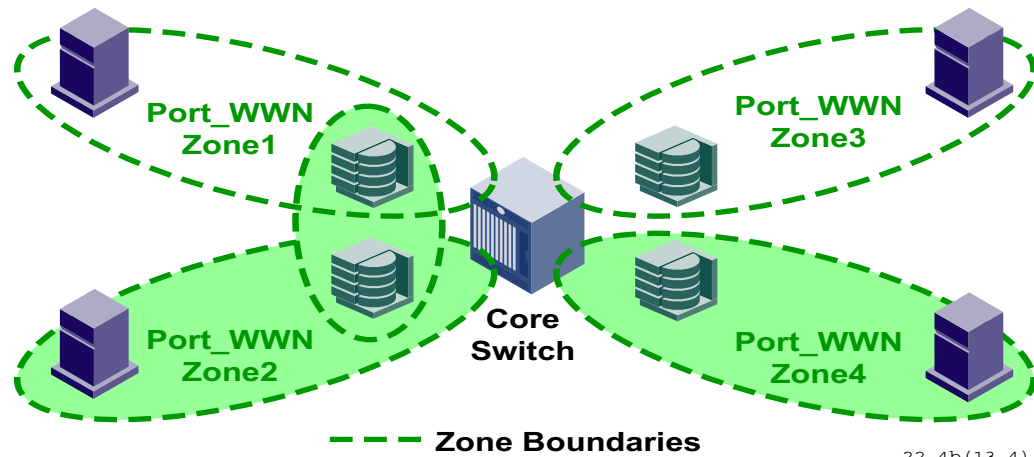
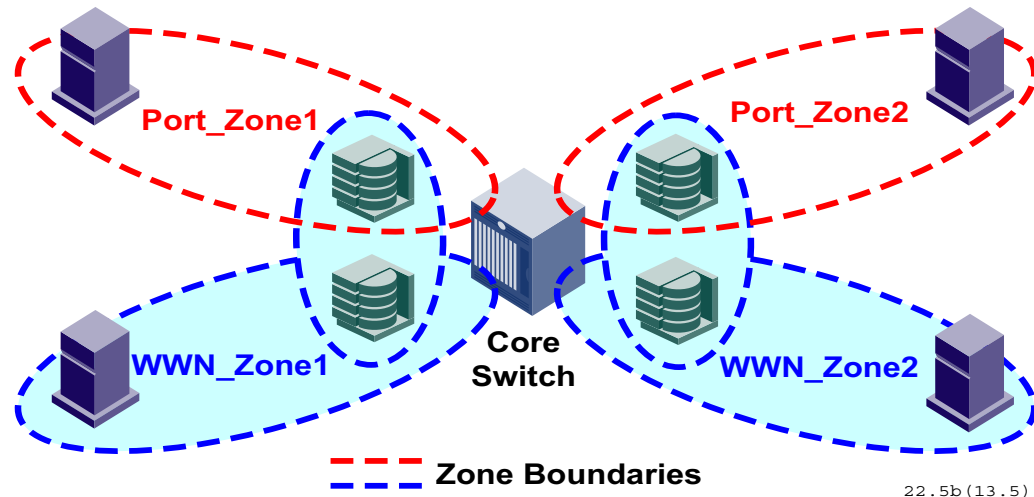


Figure 13-3 shows the same fabric components zoned in an overlapping fashion.

Figure 13-3 Hardware-Enforced Overlapping Zones



Any zone using both WWNs and *domain*, *port* entries on the Brocade 2 Gbit/sec platform relies on Name Server authentication as well as hardware-assisted (ASIC) authentication, which ensures that any PLOGI/ADISC/PDISC/ACC from an unauthorized device attempting to access a device it is not zoned with is rejected. Brocade 2 Gbit/sec switches always deploy the hardware assist, in any zone configuration (see Figure 13-4 and Figure 13-5).

Figure 13-4 Zoning with Hardware Assist (Mixed-Port and WWN Zones)**Figure 13-5** Session-Based Hard Zoning

In [Figure 13-5](#), only the ports that are overlapped are software-enforced with hardware assist.

Rules for Configuring Zones

Observe the following rules when configuring zones.

- If security is a priority, you should use hard zoning.
- The use of aliases is optional with zoning, and using aliases requires structure when defining zones. However, aliases aid administrators of a zoned fabric to understand the structure and context.
- Evaluate the security requirements of the fabric. If additional security is required, add Brocade Secure Fabric OS into the fabric.
- If the fabric includes a Brocade switch and you support a third-party switch product, they are only able to use WWN zoning; other types of zoning, including QuickLoop, are not supported.
- QuickLoop

Evaluate whether the fabric will also use QuickLoop Fabric Assist (QLFA) or QuickLoop (QL). If you are running Brocade Fabric OS v4.x, consider the following before creating and setting up QLFA zones:

- **QuickLoop Zoning.** QuickLoop/QuickLoop zones cannot run on switches running Brocade Fabric OS v4.x. However, Brocade Fabric OS v4.x can still manage (create, remove, update) QuickLoop zones on any non-v4.x switch.
- **QuickLoop Fabric Assist.** Brocade Fabric OS v4.x cannot have a Fabric Assist host directly connected to it. However, targets on a Brocade Fabric OS v4.x switch can still be part of a Fabric Assist zone if a Fabric Assist host is connected to a non-v4.x switch.
- Zone changes

Zone changes in a production fabric can cause a disruption of I/O when an RSCN is generated because of the zone change and the HBA is unable to process the RSCN fast enough. Although RSCNs are a normal part of a functioning SAN, the pause in I/O might not be acceptable. For these reasons, you should perform zone changes only when the resulting behavior is predictable and acceptable. Changing HBA drivers can rectify the situation.
- Final verification

After changing or enabling a zone configuration, confirm that the nodes and storage can identify and access one another. Depending on the platform, you might need to reboot one or more nodes in the fabric with the new changes.

The zone configuration is managed on a fabric basis. Zoning can be implemented and administered from any switch in the fabric that has an Advanced Zoning license enabled. When a change in the configuration is saved, enabled, or disabled per the transactional model, it is automatically (by closing the transaction) distributed to all switches in the fabric, preventing a single point of failure for zone information



Note

Zoning commands make changes that affect the entire fabric. When executing fabric-level configuration tasks, allow time for the changes to propagate across the fabric before executing any subsequent commands. For a large fabric, you might want to wait several minutes between commands.

Creating and Managing Zone Aliases

A zone alias is a logical group of ports, WWNs, or AL_PAs. You can simplify the process of creating zones by first specifying aliases, which eliminates the need for long lists of individual zone member names.

Be aware that if you are creating a new alias using **aliCreate w**, “1,1”, and a user in another telnet session executes **cfgEnable** (or **cfgDisable**, or **cfgSave**), the other user’s transaction will abort your transaction and you will receive an error message. Creating a new alias while there is a zone merge taking place might also abort your transaction. For more details about zone merging and zone merge conflicts, refer to [“Adding a New Switch or Fabric” on page 13-25](#).

To create an alias

1. Connect to the switch and log in as admin.
2. Enter the **aliCreate** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

Example

```
switch:admin> alicreate "array1", "2,32; 2,33; 2,34; 4,4"
switch:admin> alicreate "array2", "21:00:00:20:37:0c:66:23; 4,3"
switch:admin> alicreate "loop1", "4,6"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

To add members to an alias

1. Connect to the switch and log in as admin.
2. Enter the **aliAdd** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

Example

```
switch:admin> aliadd "array1", "1,2"
switch:admin> aliadd "array2", "21:00:00:20:37:0c:72:51"
switch:admin> aliadd "loop1", "4,6"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

To remove members from an alias

1. Connect to the switch and log in as admin.
2. Enter the **aliRemove** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

Example

```
switch:admin> aliremove "array1", "1,2"
switch:admin> aliremove "array2", "21:00:00:20:37:0c:72:51"
switch:admin> aliremove "loop1", "4,6"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

**Note**

For Fabric OS versions earlier than v4.4.0, when using the **aliRemove** command, the order in which the members appear in the list is critical. For more information on this command, refer to the *Fabric OS Command Reference Manual*.

To delete an alias

1. Connect to the switch and log in as admin.
2. Enter the **aliDelete** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

Example

```
switch:admin> alidelete "array1"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

To view an alias in the defined configuration

1. Connect to the switch and log in as admin.
2. Enter the **aliShow** command.

The following example shows all zone aliases beginning with “arr”.

Example

```
switch:admin> alishow "arr*"
alias: array1  21:00:00:20:37:0c:76:8c
alias: array2  21:00:00:20:37:0c:66:23
```

If no parameters are specified, the entire zone database (both the defined and effective configuration) is displayed.

Creating and Maintaining Zones

Before executing **cfgDisable**, **cfgEnable**, or **cfgSave**, execute the **rcsDisabled** command to check whether or not your fabric has Reliable Commit Service (RCS) enabled (**rcsDisabled=0**). If RCS is disabled (**rcsDisabled=1**), check for older switches in the fabric. After the older switches are upgraded, RCS is enabled by default.

RCS is available on all switch versions 4.1 and later. RCS guarantees that either *all or none* of the switches receive the new zone configuration. It is recommended that you use RCS to secure a reliable propagation of the latest zone configuration. If you use non-RCS mode, you must log in to every switch to monitor the status of the zone configuration.

To create a zone

1. Connect to the switch and log in as admin.
2. Enter the **zoneCreate** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

Example

```
switch:admin> zonecreate "greenzone", "2,32; 2,33; 2,34; 4,4"
switch:admin> zonecreate "redzone", "21:00:00:20:37:0c:66:23; 4,3"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

To add devices (members) to a zone

1. Connect to the switch and log in as admin.
2. Enter the **zoneAdd** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

Example

```
switch:admin> zoneadd "greenzone", "1,2"
switch:admin> zoneadd "redzone", "21:00:00:20:37:0c:72:51"
switch:admin> zoneadd "bluezone", "4,6; 21:00:00:20:37:0c:66:23"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

To remove devices (members) from a zone

1. Connect to the switch and log in as admin.
2. Enter the **zoneRemove** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

Example

```
switch:admin> zoneremove "greenzone", "1,2"
switch:admin> zoneremove "redzone", "21:00:00:20:37:0c:72:51"
switch:admin> zoneremove "bluezone", "4,6; 21:00:00:20:37:0c:66:23"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

To delete a zone

1. Connect to the switch and log in as admin.
2. Enter the **zoneDelete** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

Example

```
switch:admin> zonedeleter "bluezone"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

To view a zone in the defined configuration

1. Connect to the switch and log in as admin.
2. Enter the **zoneShow** command.

The following example shows all zones beginning with A, B, or C:

```
switch:admin> zoneshow "[A-C]*"
zone: Blue_zone 1,1; array1; 1,2; array2
zone: Bobs_zone 4,5; 4,6; 4,7; 4,8; 4,9
```

If no parameters are specified, the entire zone database (both the defined and effective configuration) is displayed.

Merging Zones

Before linking two switches together, it is important that you know the zone database limit of adjacent switches. For example, when switches running FOS v3.2, v4.4.0, or v5.x discover that the zone merge database is larger than its pre-determined zone database size limit, they issue a reject notification before symmetrically segmenting their own ends of the ISL, thereby preventing the new switch from joining the fabric.

Symmetrical segmentation occurs when both ends of an ISL are shut down. Subsequently, no frames are exchanged between those two switches.

Asymmetrical segmentation not only prevents frames from being exchanged between switches, but also causes routing inconsistencies.

The best way to avoid either type of segmentation is to know the zone database size limit of adjacent switches. The following tables provide the expected behavior based on different database sizes after a zone merge is specified.

Table 13-4 Resulting Database Size: 0 to 96K

Receiver Initiator	FOS v2.6	FOS v3.1	FOS v3.2	FOS v4.0/ v4.1/v4.2	FOS v4.3/ v4.4.0	FOS v5.0.0/ v5.0.1	Fibre Channel Router	XPath v7.3
FOS v2.6/v3.1	Join	Join	Join	Join	Join	Join	Join	Join
FOS v3.2	Join	Join	Join	Join	Join	Join	Join	Join
FOS v4.0/v4.1/ v4.2	Join	Join	Join	Join	Join	Join	Join	Join
FOS v4.3/ v4.4.0	Join	Join	Join	Join	Join	Join	Join	Join
FOS v5.0.0/ v5.0.1	Join	Join	Join	Join	Join	Join	Join	Join
Fibre Channel Router	Join	Join	Join	Join	Join	Join	Join	Join
XPath v7.3	Join	Join	Join	Join	Join	Join	Join	Join

Table 13-5 Resulting Database Size: 96K to 128K

Receiver Initiator	FOS v2.6	FOS v3.1	FOS v3.2	FOS v4.0/ v4.1/v4.2	FOS v4.3/ v4.4.0	FOS v5.0.0/ v5.0.1	Fibre Channel Router	XPath v7.3
FOS v2.6/v3.1	Segment	Segment	Segment	Segment	Segment	Segment	Join	Segment
FOS v3.2	Segment	Segment	Join	Join	Join	Join	Join	Join
FOS v4.0/v4.1/ v4.2	Segment	Segment	Segment	Join	Join	Join	Join	Join
FOS v4.3/ v4.4.0	Segment	Segment	Join	Join	Join	Join	Join	Join
FOS v5.0.0/ v5.0.1	Segment	Segment	Join	Join	Join	Join	Join	Join
Fibre Channel Router	Join	Join	Join	Join	Join	Join	Join	Join
XPath v7.3	Segment	Segment	Segment	Join	Join	Join	Join	Join

Table 13-6 Resulting Database Size: 128K to 256K

Receiver Initiator	FOS v2.6	FOS v3.1	FOS v3.2	FOS v4.0/ v4.1/v4.2	FOS v4.3/ v4.4.0	FOS v5.0.0/ v5.0.1	Fibre Channel Router	XPath v7.3
FOS v2.6/v3.1	Segment	Segment	Segment	Segment	Segment	Segment	Join	Segment
FOS v3.2	Segment	Segment	Join	Segment	Join	Join	Join	Segment
FOS v4.0/v4.1/ v4.2	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment
FOS v4.3/ v4.4.0	Segment	Segment	Join	Segment	Join	Join	Join	Segment
FOS v5.0.0/ v5.0.1	Segment	Segment	Join	Segment	Join	Join	Join	Segment
Fibre Channel Router	Join	Join	Join	Segment	Join	Join	Join	Segment
XPath v7.3	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment

Table 13-7 Resulting Database Size: 256K to 1M

Receiver Initiator	FOS v2.6	FOS v3.1	FOS v3.2	FOS v4.0/ v4.1/v4.2	FOS v4.3/ v4.4.0	FOS v5.0.0/ v5.0.1	Fibre Channel Router	XPath v7.3
FOS v2.6/v3.1	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment
FOS v3.2	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment
FOS v4.0/v4.1/ v4.2	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment
FOS v4.3/ v4.4.0	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment
FOS v5.0.0/ v5.0.1	Segment	Segment	Segment	Asymmetrical Segment	Segment	Join	Join	Segment
Fibre Channel Router	Segment	Segment	Segment	Segment	Segment	Join	Join	Segment
XPath v7.3	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment

Creating and Modifying Zoning Configurations

You can store a number of zones in a zoning configuration database. The maximum number of items that can be stored in the zoning configuration database depends on the following criteria:

- Number of switches in the fabric.
- Whether or not interoperability mode is enabled.
- Number of bytes per item. The number of bytes required for an item depends on the specifics of the fabric, but cannot exceed 64 bytes per item.

When enabling a new zone configuration, you must ensure that the size of the configuration does not exceed the minimum size supported by all switches in the fabric. This is particularly important if and when you downgrade to a FOS that supports a smaller zone database than the current FOS. In this scenario, the zone database in the current FOS would have to be changed to the smaller zone database before the downgrade.

You can use the **cfgSize** command to check both the maximum available size and the currently saved size on all switches. Refer to the *Fabric OS Command Reference Manual* for details on the **cfgSize** command. If you believe you are approaching the maximum, you can save a partially completed zoning configuration and use the **cfgSize** command to determine the remaining space.

For important considerations for managing zoning in a fabric, and more details about the maximum zone database size for each version of the FOS, refer to [“Maintaining Zone Objects” on page 13-22](#).

To create a zoning configuration

1. Connect to the switch and log in as admin.
2. Enter the **cfgCreate** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

Example

```
switch:admin> cfgcreate "NEW_cfg", "redzone; bluezone; greenzone"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

To add zones (members) to a zoning configuration

1. Connect to the switch and log in as admin.
2. Enter the **cfgAdd** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

Example

```
switch:admin> cfgadd "newcfg", "bluezone"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

To remove zones (members) from a zone configuration

1. Connect to the switch and log in as admin.
2. Enter the **cfgRemove** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

Example

```
switch:admin> cfgremove "newcfg", "redzone"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

To delete a zone configuration

1. Connect to the switch and log in as admin.
2. Enter the **cfgDelete** command:
3. Enter the **cfgSave** command to save the change to the defined configuration.

Example

```
switch:admin> cfgdelete "testcfg"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

To clear changes to a configuration

Use the **cfgTransAbort** command. When this command is executed, all changes since the last save operation (performed with the **cfgSave** command) are cleared.

In the following example, assume that the removal of a member from **zone1** was done in error:

```
switch:admin> zoneremove "zone1","3,5"
switch:admin> cfgtransabort
```

To view all zone configuration information

If you do not specify an operand when executing the **cfgShow** command to view zone configurations, then all zone configuration information (both defined and effective) displays. If there is an outstanding transaction, then the newly edited zone configuration that has not yet been saved is displayed. If there are no outstanding transactions, then the committed zone configuration displays.

1. Connect to the switch and log in as admin.
2. Enter the **cfgShow** command with no operands.

Example

```
switch:admin> cfgshow
Defined configuration:
  cfg:   USA1      Blue_zone
  cfg:   USA_cfg Red_zone; Blue_zone
  zone:  Blue_zone
        1,1; array1; 1,2; array2
  zone:  Red_zone
        1,0; loop1
  alias: array1  21:00:00:20:37:0c:76:8c; 21:00:00:20:37:0c:71:02
  alias: array2  21:00:00:20:37:0c:76:22; 21:00:00:20:37:0c:76:28
  alias: loop1   21:00:00:20:37:0c:76:85; 21:00:00:20:37:0c:71:df

Effective configuration:
  cfg:   USA_cfg
  zone:  Blue_zone
        1,1
        21:00:00:20:37:0c:76:8c
        21:00:00:20:37:0c:71:02
        1,2
        21:00:00:20:37:0c:76:22
        21:00:00:20:37:0c:76:28
  zone:  Red_zone
        1,0
        21:00:00:20:37:0c:76:85
        21:00:00:20:37:0c:71:df
```

To view selected zone configuration information

1. Connect to the switch and log in as admin.
2. Enter the **cfgShow** command and specify a pattern.

For example, to display all zone configurations that start with “Test”:

```
switch:admin> cfgshow "Test*"
cfg:   Test1 Blue_zone
cfg:   Test_cfg Red_zone; Blue_zone
```

To view a configuration in the effective zone database

1. Connect to the switch and log in as admin.
2. Enter the **cfgActvShow** command.

Example

```
switch:admin> cfgactvshow
Effective configuration:
  cfg:   NEW_cfg
  zone:  Blue_zone
    1,1
    21:00:00:20:37:0c:76:8c
    21:00:00:20:37:0c:71:02
    1,2
    21:00:00:20:37:0c:76:22
    21:00:00:20:37:0c:76:28
  zone:  Red_zone
    1,0
    21:00:00:20:37:0c:76:85
    21:00:00:20:37:0c:71:df
```

Maintaining Zone Objects

While you can use the **cfgDelete** command to delete a zone configuration, there is a quicker and easier way to perform the same task via the zone object commands (**zoneObjectExpunge**, **zoneObjectCopy**, and **zoneObjectRename**). You can also copy and rename zone objects. When you copy a zone object, the resulting object has the same type as the original. Deleting a zone object also removes the object from any member lists of other objects. You can rename objects for all zone object types.

To copy a zone object

1. Connect to the switch and log in as admin.
2. Enter the **cfgShow** command to view the zone configuration objects you want to copy.

For example, to display all zone configuration objects that start with “Test”:

```
switch:admin> cfgshow "Test*"
cfg:   Test1 Blue_zone
cfg:   Test_cfg Red_zone; Blue_zone
```

3. Enter the **zoneObjectCopy** command, specifying the zone configuration objects you want to copy, along with the new object name. Note that zone configuration names are case sensitive; blank spaces are ignored.

```
switch:admin> zoneobjectcopy "Test1", "US_Test1"
```

4. Enter the **cfgShow** command to verify the new zone object is present.

```
switch:admin> cfgshow "Test*"
cfg:   Test1 Blue_zone
cfg:   Test_cfg Red_zone; Blue_zone
cfg:   US_Test1 Blue_zone
```

5. If you want the change preserved when the switch reboots, save it to nonvolatile (also known as “flash”) memory by entering the **cfgSave** command.
6. For the change to become effective, enable the appropriate zone configuration using the **cfgEnable** command.

To delete a zone object

1. Connect to the switch and log in as admin.
2. Enter the **cfgShow** command to view the zone configuration objects you want to delete.

```
switch:admin> cfgShow
Defined configuration:
cfg: USA_cfg Red_zone; White_zone; Blue_zone
zone: Blue_zone
      1,1; array1; 1,2; array2
zone: Red_zone
      1,0; loop1
zone: White_zone
      1,3; 1,4
alias: array1 21:00:00:20:37:0c:76:8c; 21:00:00:20:37:0c:71:02
alias: array2 21:00:00:20:37:0c:76:22; 21:00:00:20:37:0c:76:28
alias: loop1 21:00:00:20:37:0c:76:85; 21:00:00:20:37:0c:71:df
```

3. Enter the **zoneObjectExpunge** command to delete the zone object. Note that zone configuration names are case sensitive; blank spaces are ignored

```
switch:admin> zoneObjectExpunge "White_zone"
```

4. Enter the **cfgShow** command to verify the deleted zone object is no longer present.
5. If you want the change preserved when the switch reboots, save it to nonvolatile (also known as “flash”) memory by entering the **cfgSave** command.
6. For the change to become effective, enable the appropriate zone configuration using the **cfgEnable** command.

To rename a zone object

1. Connect to the switch and log in as admin.
2. Enter the **cfgShow** command to view the zone configuration objects you want to rename.

```
switch:admin> cfgShow
Defined configuration:
cfg: USA_cfg Red_zone; White_zone; Blue_zone
zone: Blue_zone
      1,1; array1; 1,2; array2
zone: Red_zone
      1,0; loop1
zone: White_zone
      1,3; 1,4
alias: array1 21:00:00:20:37:0c:76:8c; 21:00:00:20:37:0c:71:02
alias: array2 21:00:00:20:37:0c:76:22; 21:00:00:20:37:0c:76:28
alias: loop1 21:00:00:20:37:0c:76:85; 21:00:00:20:37:0c:71:df
```

3. Enter the **zoneObjectRename** command to rename zone configuration objects. Note that zone configuration names are case sensitive; blank spaces are ignored

```
switch:admin> zoneObjectRename "White_zone", "Red_zone"
```

4. Enter the **cfgShow** command to verify the renamed zone object is present.
5. If you want the change preserved when the switch reboots, save it to nonvolatile (also known as “flash”) memory by entering the **cfgSave** command.
6. For the change to become effective, enable the appropriate zone configuration using the **cfgEnable** command.

For more details about the **zoneObjectCopy**, **cfgShow**, **cfgEnable**, and **cfgSave** commands, refer to the *Fabric OS Command Reference Manual*.

Managing Zoning Configurations in a Fabric

To modify an existing zone configuration, you can add, delete, or remove individual elements to create the desired configuration. After the changes have been made, save the configuration to ensure the configuration is permanently saved in the switch and that the configuration is replicated throughout the fabric.

The switch configuration file can also be uploaded to the host for archiving and it can be downloaded from the host to a switch in the fabric. Refer to [“Backing Up a Configuration” on page 4-1](#), [“Restoring a Configuration” on page 4-2](#), or the **configUpload** and **configDownload** commands in the *Fabric OS Command Reference Manual*.

[Table 13-8](#) presents zoning database size limitations for various Fabric OS release versions. The maximum size of a zone database is the upper limit for the defined configuration, and it is determined by the amount of flash memory available for storing the defined configuration.

Table 13-8 Zoning Database Limitations

Fabric OS Version	Maximum Database Size (KB))
2.4.0	64
2.5.0	64
2.6.0	96
3.0.0	128
3.1.0	96
3.2.0	256
4.0.0, 4.1.0, 4.2.0	128
4.4.0	256
5.0.1	256

Adding a New Switch or Fabric

When a new switch is added to the fabric, it automatically takes on the zone configuration information from the fabric. Use the **cfgActvShow** command to verify that the zoning information is the same on each switch in the fabric.

If you are adding a switch that is already configured for zoning, use the **cfgClear** and **cfgSave** commands (or use **cfgClear** and **cfgDisable** if there is an effective configuration) before connecting it to the zoned fabric.

Adding a new fabric that has no zone configuration information to an existing fabric is very similar to adding a new switch. All switches in the new fabric inherit the zoning configuration data. If a zone configuration is in effect, then the same configuration becomes the enabled configuration. The **cfgActvShow** command will display the same information on all switches in the newly formed fabric.

Before the new fabric can merge successfully, it must pass the following criteria:

- Before merging zones

To facilitate merging, check the following before merging switches or fabrics.

- **Zoning licenses**—All switches must have a Zoning license enabled.
- **Native operating mode**—All switches must be in the native operating mode.
- **Brocade Secure Fabric OS**—If one switch has Brocade Secure Fabric OS enabled, all switches in the fabric must have Brocade Secure Fabric OS. Refer to the *Secure Fabric OS Manual* for more information.

- Merging and segmentation

The fabric is checked for segmentation during power-up or when a switch is disabled or enabled, or when a new switch is added.

The database is the zone configuration database. This is the data displayed as the “defined configuration” in the **cfgShow** command. It is stored in nonvolatile memory by the **cfgSave** command. This database is a replicated database, which means that all switches in the fabric will have copy of this database. When a change is made to the defined configuration, the switch where the changes were made must close its transaction for the change to get propagated throughout the fabric.

- Merging rules

Observe these rules when merging zones:

Local and adjacent configurations	If the local and adjacent zone database configurations are the same, they will remain unchanged after the merge.
Effective configurations	If there is an effective configuration between two switches, the zone configuration in effect match.
Zone object naming	If a zoning object has the same name in both the local and adjacent defined configurations, the object types and member lists must match. When comparing member lists, the content and order of the members are important.

Objects in adjacent configurations	If a zoning object appears in an adjacent defined configuration, but not in the local defined configuration, the zoning object is added to the local defined configuration. The modified zone database must fit in the nonvolatile memory area allotted for the zone database.
Local configuration modification	If a local defined configuration is modified because of a merge, the new zone database is propagated to other the switches within the merge request.

- Merging Two Fabrics

Both fabrics have identical zones and configurations enabled. The two fabrics will join to make one larger fabric with the same zone configuration across the newly created fabric.

If the two fabrics have different zoning configurations, they will be merged. If the two fabrics cannot join, the ISL between the switches will be segmented.

- Merge Conflicts

When a merge conflict is present, a merge will not take place and the ISL will segment. Use the **switchShow** or **errLogShow** commands to obtain additional information about possible merge conflicts, because many non-zone related configuration parameters can cause conflicts

If the fabrics have different zone configuration data, the system attempts to merge the two sets of zone configuration data. If the zones cannot merge, the ISL will be segmented.

A merge is not possible if any of the following conditions exist:

Configuration mismatch	Zoning is enabled in both fabrics and the zone configurations that are enabled are different in each fabric.
Type mismatch	The name of a zone object in one fabric is used for a different type of zone object in the other fabric.
Content mismatch	The definition of a zone object in one fabric is different from the definition of zone object with the same name in the other fabric.



Note

If the zoneset members on two switches are not listed in the same order, the configuration is considered a mismatch, resulting in the switches being segmented from the fabric. For example: `cfg1 = z1; z2` is different from `cfg1 = z2; z1`, even though members of the configuration are the same. If zoneset members on two switches have the same names defined in the configuration, make sure zoneset members are listed in the same order.

Splitting a Fabric

If the connections between two fabrics are no longer available, the fabric will segment into two separate fabrics. Each new fabric will retain the same zone configuration.

If the connections between two fabrics are replaced and no changes have been made to the zone configuration in either of the two fabrics, then the two fabrics will merge back into one single fabric. If any changes that cause a conflict have been made to either zone configuration, then the fabrics might segment.

Using Zoning to Administer Security

Zones provide controlled access to fabric segments and establish barriers between operating environments. They isolate systems with different uses, protecting individual systems in a heterogeneous environment; for example, when zoning is in secure mode, no merge operations occur.

Brocade Advanced Zoning is configured on the primary Fabric Configuration Server (FCS). The primary FCS switch makes zoning changes and other security-related changes. The primary FCS switch also distributes zoning to all other switches in the secure fabric. All existing interfaces can be used to administer zoning (depending on the policies; refer to the *Secure Fabric OS Administrator's Guide* for information about security policies).

You must perform zone management operations from the primary FCS switch using a zone management interface, such as telnet or Advanced Web Tools. You can alter a zoning database, provided you are connected to the primary FCS switch.

When two secure fabrics join, the traditional zoning merge does not occur. Instead, a zoning database is downloaded from the primary FCS switch of the merged secure fabric. When E_Ports are active between two switches, the name of the FCS server and a zoning policy set version identifier are exchanged between the switches. If the views of the two secure fabrics are the same, the fabric's primary FCS server downloads the zoning database and security policy sets to each switch in the fabric. If there is a view conflict, the E_Ports are segmented due to incompatible security data.

As part of zoning architecture, you must determine which of the two basic zoning architectures (hard or soft) works best for your fabric. With time and planning, the basic hard zone configuration works for most sites. If a site has additional security needs, use the additional layer of Secure Fabric OS, apart from the standard zoning architecture.



Note

Secure Fabric OS requires the activation of a Brocade security license and an Advanced Zoning license.

Resolving Zone Conflicts

Zone conflicts can be resolved by saving a configuration file with the **configUpload** command, examining the zoning information in the file, and performing a cut and paste operation so that the configuration information matches in the fabrics being merged.

After examining the configuration file, you can choose to resolve zone conflicts by using the **cfgClear** command followed by the **cfgDisable** command on the incorrectly configured segmented fabric, followed by a **portDisable/portEnable** command on one of the ISL ports that connects the fabrics. This will cause a merge, making the fabric consistent with the correct configuration.



Caution

Be careful using the **cfgClear** command, because it deletes the defined configuration.

Table 13-9 lists considerations for zoning architecture.

Table 13-9 Considerations for Zoning Architecture

Item	Description
Type of zoning: hard or soft (session-based)	If security is a priority, hard zoning is recommended.
Use of aliases	The use of aliases is optional with zoning. Using aliases requires structure when defining zones. Aliases will aid administrators of zoned fabric in understanding the structure and context.
Security requirements	Evaluate the security requirements of the fabric. If additional security is required, add Brocade Secure Fabric OS into the fabric.
Interoperability Fabric	If the fabric includes a third-party switch product, only WWN zoning is supported. Other types of zoning, including QuickLoop, are not supported.
QLFA zones	<p>Evaluate if the fabric will have QuickLoop Fabric Assist (QLFA) or QuickLoop (QL) in it, and consider the following items before creating and setting up QLFA zones:</p> <p>QuickLoop Zoning—QuickLoop/QuickLoop zones cannot run on Fabric OS v4.1.0 or later. However, Fabric OS can manage (create, remove, update) QL zones.</p> <p>QuickLoop Fabric Assist—A switch running Fabric OS v4.1.0 or later cannot have a Fabric Assist host directly connected to it. However, such a switch can be part of a Fabric Assist zone if a Fabric Assist host is connected to a compatible switch in the fabric.</p>
Testing	Testing a (new) zone configuration. Before implementing a zone, the user should run the Zone Analyzer from Web Tools to isolate any possible problems. This is especially useful as fabrics increase in size.

Table 13-9 Considerations for Zoning Architecture (Continued)

Item	Description
Effect of changes in a production fabric	Zone changes in a production fabric can result in a disruption of I/O under conditions where an RSCN is issued as a result of a zone change and the HBA is unable to process the RSCN fast enough. Though RSCNs are a normal part of a functioning SAN, the pause in I/O might not be acceptable. For these reasons, it is recommended to perform zone changes only when the resulting behavior is predictable and acceptable. Changing HBA drivers can rectify the situation.
Confirming operation	After changing or enabling a zone configuration, the user should confirm that the nodes and storage are able to identify and access one another. Depending on the platform, the user might need to reboot one or more nodes in the fabric with the new changes.

Administering Advanced Performance Monitoring

This topic contains procedures for the Brocade Advanced Performance Monitoring licensed feature:

Based on Brocade Frame Filtering technology and a unique performance counter engine, Advanced Performance Monitoring is a comprehensive tool for monitoring the performance of networked storage resources. It supports direct-attach, loop, and switched fabric Fibre Channel SAN topologies by:

- Monitoring transaction performance from source to destination.
- Reporting cyclic redundancy check (CRC) error measurement statistics.
- Measuring Brocade InterSwitch Link (ISL) Trunking performance and resource usage.

Further features are provided through Web Tools:

- Measuring device performance by port, arbitrated loop physical address (AL_PA), and logical unit number (LUN).
- Comparing IP versus SCSI traffic on each port.
- Providing a library of predefined graphs.



Note

The SilkWorm 200E switch running Fabric OS v5.0.1 does not display AL_PA measurements for end-to-end monitors. It provides port CRC reports through Web Tools.

The SilkWorm 4100 switch running Fabric OS v4.4.0 does not display AL_PA measurements for end-to-end monitors. It provides port CRC reports through Web Tools.

Table 14-1 lists commands associated with Advanced Performance Monitoring. For detailed information on these commands, refer to the *Fabric OS Command Reference Manual*.

Table 14-1 Advanced Performance Monitoring Commands

Command	Description
perfAddEEMonitor	Add an end-to-end monitor to a port.
perfAddIPMonitor	Add an IP monitor to a port.
perfAddReadMonitor	Add a SCSI Read monitor to a port.
perfAddRwMonitor	Add a SCSI Read and Write monitor to a port.
perfAddSCSIMonitor	Add a SCSI traffic frame monitor to a port.
perfAddUserMonitor	Add a filter-based monitor to a port.
perfAddWriteMonitor	Add a SCSI Write monitor to a port.
perfCfgClear	Clear the performance monitoring settings from nonvolatile (flash) memory.
perfCfgRestore	Restore performance monitoring settings from nonvolatile (flash) memory.
perfCfgSave	Save the current performance monitoring settings to nonvolatile (flash) memory.
perfClrAlpaCrc	Clear an AL_PA device CRC count by the port and AL_PA.
perfClearEEMonitor	Clear end-to-end monitor counters on a port.
perfClearFilterMonitor	Clears filter-based monitor counters.
perfDelEEMonitor	Delete an end-to-end monitor on port.
perfDelFilterMonitor	Delete a filter-based monitor.
perfMonitorClear	Clear statistics counters of end-to-end, filter-based, and ISL monitors on a port.
perfMonitorShow	Display end-to-end, filter-based, and ISL monitors on a port.
perfSetPortEEMask	Set overall mask for end-to-end (EE) monitors.
perfShowAlpaCrc	Display the AL_PA CRC count by port or by AL_PA.
perfShowEEMonitor	Show user-defined end-to-end monitors.
perfShowFilterMonitor	Show filter-based monitors.
perfShowPortEEMask	Display the current end-to-end mask of a port.



Note

The command examples in this chapter use the slot/port syntax required by SilkWorm 12000, 24000, and 48000 directors. For SilkWorm 200E, 3016, 3250, 3850, 3900, 4012, and 4100 switches, use only the port number where needed in the commands.

Displaying and Clearing the CRC Error Count

You can use the **perfShowAlpaCrc** command to display the CRC error count for all AL_PA devices or for a single AL_PA on a specific active L_Port.

Example: Displaying the CRC error count for all AL_PA devices on a port

```
switch:admin> perfshowalpacrc 1/1
AL_PA      CRC count
-----
0xd9              0
```

Example: Displaying the CRC error count for a single AL_PA device on a port

```
switch:admin> perfshowalpacrc 1/1, 0xd9
The CRC count at ALPA 0xd9 on port 1 is 0x00000000.
```

Example: Clearing the CRC error count

```
switch:admin> perfclrallpacrc 1/1, 0xd9
CRC error count at AL_PA 0xd9 on port 1 is cleared.
switch:admin> perfclrallpacrc 1/1
No AL_PA value is specified. This will clear all AL_PA CRC
counts on port 1. Do you want to continue? (yes, y, no, n): [no] y
Please wait ...
All alpa CRC counts are cleared on port 1.
```

In v3.1.0, v4.1.0, and later versions, you can use the **portStatsClear** command clears AL_PA- based CRC error counters for all the ports in the same group.

Monitoring End-to-End Performance

End-to-end performance monitoring counts the number of words and CRC errors in Fibre Channel frames for a specified Source ID (SID) and Destination ID (DID) pair. An end-to-end performance monitor includes these counts:

- RX_COUNT (words in frames received at the port)
- TX_COUNT (words in frames transmitted from the port)
- CRC_COUNT (frames with CRC errors received at or transmitted from the port)

To enable end-to-end performance monitoring, you must configure an end-to-end monitor on a port, specifying the SID-DID pair (in hexadecimal). The monitor counts only those frames with matching SID and DID.

Each SID or DID has three fields, listed in the following order:

- Domain ID (DD)
- Area ID (AA)
- AL_PA (PP)

For example, the SID 0x118a0f denotes DD 0x11, AA 0x8a, and AL_PA 0x0f.

You can monitor end-to-end performance using the **perfMonitorShow** command, as described in [“Displaying Monitor Counters” on page 14-11](#). You can clear end-to-end counters using the **perfMonitorClear** command, as described in [“Clearing Monitor Counters” on page 14-14](#).



Note

For end-to-end monitors, CRC counters are not displayed on the SilkWorm 200E, 4012, 4100, and SilkWorm 48000 switches.

Adding End-to-End Monitors

An end-to-end monitor counts the following items for a port: number of words received, number of words transmitted, and number of CRC errors detected in frames.

SilkWorm 200E, 3016, 3250, 3850, 3900, 4012, 12000, and 24000 models allow up to eight end-to-end monitors.

The SilkWorm 4100 and 48000 models allow up to 256 end-to-end monitors shared by all ports. (The number of interswitch links configured on the switch affects the amount of resources available for end-to-end monitors.)

End-to-end monitors cannot be added to interswitch links.

The monitor count is qualified using either of following conditions:

- For frames received at the port with the end-to-end monitor installed, the frame SID is the same as “SourceID” and the frame DID is the same as “DestID”. The RX_COUNT and CRC_COUNT are updated accordingly.
- For frames transmitted from the port with the end-to-end monitor installed, the frame DID is the same as “SourceID” and the frame SID is the same as “DestID”. The TX_COUNT and CRC_COUNT are updated accordingly.

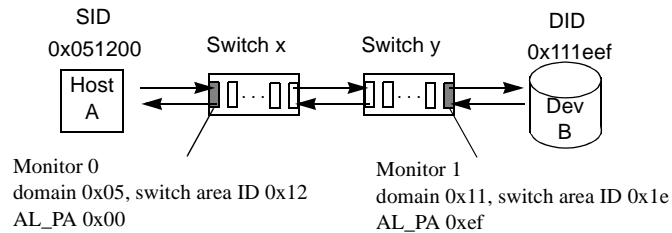


Note

How the area ID for a port relates to the port number depends upon the PID format used by the fabric. Refer to [“Configuring the PID Format” on page A-1](#) for more information.

[Figure 14-1](#) shows two devices:

- Host A is connected to domain 5 (0x05), switch area ID 18 (0x12), AL_PA 0x00 on Switch X
- Dev B is a storage device connected to domain 17 (0x11), switch area ID 30 (0x1e), AL_PA 0xef on Switch Y.

Figure 14-1 Setting End-to-End Monitors on a Port**Note**

End-to-end performance monitoring looks at traffic on the receiving port respective to the SID only. In [Figure 14-1](#), if you add a monitor to slot 2, port 2 on Switch x, specifying Dev B as the SID and Host A as the DID, no counters (except CRC) will be incremented.

To monitor the traffic from Host A to Dev B

```
switch:admin> perfaddeemonitor 2/2, "0x051200" "0x111eef"
End-to-End monitor number 0 added.
```

Add Monitor 0 to slot 2, port 2 on Switch x, specifying 0x051200 as the SID and 0x111eef as the DID, as shown in the following example:

Monitor 0 counts the frames that have an SID of 0x051200 and a DID of 0x111eef. For monitor 0, RX_COUNT is the number of words from Host A to Dev B, TX_COUNT is the number of words from Dev B to Host A, and CRC_COUNT is the number of frames in both directions with CRC errors.

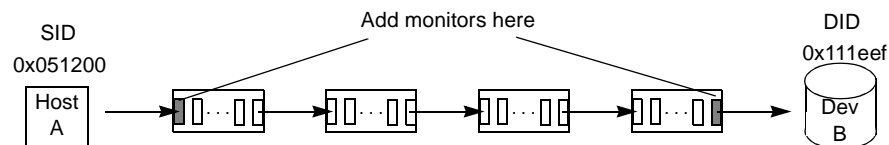
To monitor the traffic from Dev B to Host A

```
switch:admin> perfaddeemonitor 2/14, "0x111eef" "0x051200"
End-to-End monitor number 1 added.
```

Add Monitor 1 to slot 2, port 14 on Switch y, specifying 0x111eef as the SID and 0x051200 as the DID, as shown in the following example.

Monitor 1 counts the frames that have an SID of 0x111eef and a DID of 0x051200. For monitor 1, RX_COUNT is the number of words from Dev B to Host A, TX_COUNT is the number of words from Host A to Dev B, and CRC_COUNT is the number of frames in both directions with CRC errors.

[Figure 14-2](#) shows several switches and the correct ports on which to add performance monitors for a specified SID-DID pair.

Figure 14-2 Proper Placement of End-to-End Performance Monitors

Setting a Mask for End-to-End Monitors

End-to-end monitors count the number of words in Fibre Channel frames that match a specific SID/DID pair. If you want to match only part of the SID or DID, you can set a mask on the port to compare only certain parts of the SID or DID. By default, the frame must match the entire SID and DID to trigger the monitor. By setting a mask, you can choose to have the frame match only one or two of the three fields (Domain ID, Area ID, and AL_PA) to trigger the monitor.



Note

Only one mask per port can be set. When you set a mask, all existing end-to-end monitors are deleted.

You can specify a mask using the **perfSetPortEeMask** command in the form `dd:aa:pp`, where `dd` is the domain ID mask, `aa` is the area ID mask, and `pp` is the AL_PA mask. The values for `dd`, `aa`, and `pp` are either `ff` (the field must match) or `00` (the field is ignored). The default EE mask value is `ff:ff:ff`. The command sets the mask for all end-to-end monitors of a port. If any end-to-end monitors are programmed on a port when the **perfSetPortEeMask** command is issued, a message displays such as that in the following example:

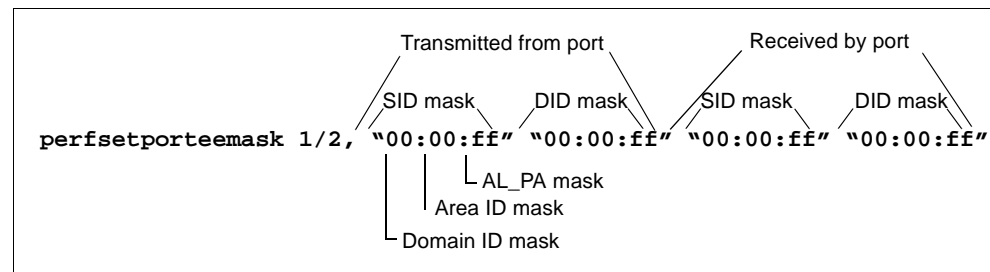
```
switch:admin> perfsetporteemask 1/2, "00:00:ff"
EE monitors are currently programmed on this port. Changing EE mask
for this port will cause ALL EE monitors on this port to be deleted.
Do you want to continue? (yes, y, no, n): [no] y

EE mask on port <port-number> is set and EE monitors were deleted
```

The **perfSetPortEeMask** command sets a mask for the Domain ID, Area ID, and AL_PA of the SIDs and DIDs for frames transmitted from and received by the port.

Figure 14-3 shows the mask positions in the command. A mask ("ff") is set on slot 1, port 2 to compare the AL_PA fields on the SID and DID in all frames (transmitted and received) on port 2. The frame SID and DID must match only the AL_PA portion of the specified SID-DID pair. Each port can have only one EE mask. The mask is applied to all end-to-end monitors on the port. Individual masks for each monitor on the port cannot be specified.

Figure 14-3 Mask Positions for End-to-End Monitors



To display the current end-to-end mask of a port

Enter the **perfShowPortEeMask** command.

The end-to-end mask has 12 fields, each with a value of on or off.

Example: Setting and displaying an end-to-end mask

```

switch:admin> perfsetporteemask 1/11, "00:00:ff" "00:00:ff" "00:00:ff" "00:00:ff"
The EE mask on port 11 is set and EE counters are reset.
switch:admin> perfshowporteemask 1/11
The EE mask on port 11 is set by application TELNET
TxSID Domain:  off
TxSID Area:    off
TxSID AL_PA:   on
TxDID Domain:  off
TxDID Area:    off
TxDID AL_PA:   on
RxSID Domain:  off
RxSID Area:    off
RxSID AL_PA:   on
RxDID Domain:  off
RxDID Area:    off
RxDID AL_PA:   on

```

The end-to-end mask is set on slot 1, port 11.

To display a monitor

Enter the **perfMonitorShow** command, as described in [“Displaying Monitor Counters” on page 14-11](#).

Deleting End-to-End Monitors

Enter the **perfDelEeMonitor** command to delete end-to-end monitors. You can delete all monitors or specific monitors. The following example deletes the end-to-end monitor number 0 on slot 1, port 2:

```

switch:admin> perfdeleemonitor 1/2, 0
End-to-End monitor number 0 deleted

```

Monitoring Filter-Based Performance

Filter-based performance monitoring counts the number of times a frame with a particular pattern is transmitted by a port. Filter-based monitoring is achieved by configuring a filter for a particular purpose. The filter can be a standard filter (for example, a SCSI read command filter that counts the number of SCSI read commands that have been transmitted by the port) or a user-defined filter customized for your particular use.

For the SilkWorm 4100 models, the maximum number of filters is 12 per port, in any combination of standard filters and user-defined filters.

For SilkWorm 200E, 3016, 3250, 3850, 3900, 4012, 12000, and 24000 models, the maximum number of filters is eight per port, in any combination of standard filters and user-defined filters.

For the SilkWorm 48000 models, the maximum number of filters is 12 per port.

The actual number of filters that can be configured on a port depends on the complexity of the filters. For trunked ports, the filter is configured on the trunk master.

You can monitor filter-based performance using the **perfMonitorShow** command, as described in “[Displaying Monitor Counters](#)” on page 14-11. You can clear filter-based counters using the **perfMonitorClear** command, as described in “[Clearing Monitor Counters](#)” on page 14-14.

Adding Standard Filter-based Monitors

Table 14-2 lists the commands for adding standard filter-based monitors to a port.

Table 14-2 Commands to Add Filter-Based Monitors

Telnet command	Description
perfAddReadMonitor	Count the number of SCSI Read commands
perfAddWriteMonitor	Count the number of SCSI Write commands
perfAddRwMonitor	Count the number of SCSI Read and Write commands
perfAddScsiMonitor	Count the number of SCSI traffic frames
perfAddIpMonitor	Count the number of IP traffic frames

The following example adds filter-based monitors to slot 1, port 2 and displays the results:

```
switch:admin> perfaddreadmonitor 1/2
SCSI Read filter monitor #0 added
switch:admin> perfaddwritemonitor 1/2
SCSI Write filter monitor #1 added
switch:admin> perfaddrwmonitor 1/2
SCSI Read/Write filter monitor #2 added
switch:admin> perfaddscsimonitor 1/2
SCSI traffic frame monitor #3 added
switch:admin> perfaddipmonitor 1/2
IP traffic frame monitor #4 added
switch:admin> perfmonitorshow --class FLT 1/2
There are 5 filter-based monitors defined on port 2.
```

KEY	ALIAS	OWNER_APP	OWNER_IP_ADDR	FRAME_COUNT
0	SCSI Read	TELNET	N/A	0x0000000000000000
1	SCSI Write	TELNET	N/A	0x0000000000000000
2	SCSI R/W	TELNET	N/A	0x0000000000000000
3	SCSI Frame	TELNET	N/A	0x0000000000000000
4	IP Frame	TELNET	N/A	0x0000000000000000

Adding Custom Filter-Based Monitors

In addition to the standard filters—read, write, read/write, SCSI frame and IP frame—you can create custom filters to gather statistics that fit your needs.

To define a custom filter, use the **perfAddUserMonitor** command. With this command, you must specify a series of *offsets*, *masks*, and *values*. For all transmitted frames, the switch performs these tasks:

- Locates the byte found in the frame at the specified *offset*.

- Applies the *mask* to the byte found in the frame.
- Compares the value with the given *values* in the **perfAddUserMonitor** command.
- Increments the filter counter if a match is found.

The following number of offsets can be specified:

- SilkWorm 3016, 3250, 3850, 3900, 12000, and 24000 models (Fabric OS v4.0.0 or later)
Up to two different offsets per port (one offset when FICON management server mode (FMS) is enabled).
- SilkWorm 48000 model (Fabric OS v5.0.1 or later)
Up to 15 different offsets per port (14 offsets when fmsmode is enabled).
- SilkWorm 3200, 3800 models (Fabric OS v3.0.0 or later)
Up to three different offsets per port.
- SilkWorm 4100 model (Fabric OS v4.4.0 or later)
Up to 15 different offsets per port (14 offsets when FMS is enabled).
- SilkWorm 200E and 4012 models (Fabric OS v5.0.1)
Up to 7 different offsets per port (6 offsets when FMS is enabled).

You can specify up to four values to compare against each offset. If more than one offset is required to properly define a filter, the bytes found at each offset must match one of the given values for the filter to increment its counter. If one or more of the given offsets does not match any of the given values, the counter does not increment. The value of the offset must be between 0 and 63, in decimal format. Byte 0 indicates the first byte of the Start of Frame (SOF), byte 4 is the first byte of the frame header, and byte 28 is the first byte of the payload. Thus only the SOF, frame header, and first 36 bytes of payload can be selected as part of a filter definition. Offset 0 is a special case, which can be used to monitor the first 4 bytes of the frame (SOF). When the offset is set to 0, the values 0–7 that are checked against that offset are predefined as shown in [Table 14-3](#).

Table 14-3 Predefined Values at Offset 0

Value	SOF	Value	SOF
0	SOFf	4	SOFi2
1	SOFc1	5	SOFn2
2	SOFi1	6	SOFi3
3	SOFn1	7	SOFn3

If the switch does not have enough resources to create a given filter, then other filters might have to be deleted to free resources.

To add filter-based monitors

```
switch:admin> perfaddusermonitor 4/2, "12, 0xff, 0x05, 0x08; 9, 0xff, 0x02" "FCP/IP"
User monitor #5 added
switch:admin> perfaddusermonitor 1/2, "0, 0xff, 6"
User Monitor #6 added
```

Two filter-based monitors are added. The first monitor (#5) counts all FCP and IP frames transmitted from domain 0x02 for slot 4, port 2. The FCP and IP protocols are selected by monitoring offset 12, mask 0xff and matching values of 0x05 or 0x08. Domain 2 is selected by monitoring offset 9, mask 0xff, and matching a value of 0x02. The monitor counter is incremented for all outgoing frames from port 2 where byte 9 is 0x02 and byte 12 is 0x05 or 0x08.

The second monitor (#6) is for SOFi3 on slot 1, port 2.

Deleting Filter-Based Monitors

To delete a filter-based monitor:

1. List the valid monitor numbers by entering the **perfShowFilterMonitor** command.
2. Enter the **perfDelFilterMonitor** command to delete a specific monitor. If you do not specify which monitor number to delete, you are asked if you want to delete all entries.

The following example displays the monitors on slot 1, port 4 using the **perfShowFilterMonitor** command (the monitor numbers are listed in the KEY column) and deletes monitor number 1 on slot 1, port 4 using the **perfDelFilterMonitor** command:

```
switch:admin> perfshowfiltermonitor 1/4
There are 4 filter-based monitors defined on port 4.
KEY    ALIAS    OWNER_APP    OWNER_IP_ADDR    FRAME_COUNT
-----
0  SCSI Read  TELNET                N/A              0x0000000000002208
1  SCSI Write TELNET                N/A              0x000000000000464a
2  SCSI R/W   TELNET                N/A              0x000000000000fd8c
3  SCSI Frame WEB_TOOLS           192.168.169.40  0x000000000002c229
switch:admin> perfdelfiltermonitor 1/4, 1
The specified filter-based monitor is deleted.
```

Monitoring ISL Performance

ISL monitoring is set up on E_Ports automatically in release v4.4.0 and later.

An ISL monitor measures traffic to all reachable destination domains for an ISL, showing which destination domain is consuming the most traffic. If there are more than 16 domains, the monitor samples traffic and extrapolates the measurement.

You can monitor ISL performance using the **perfMonitorShow** command, as described in [“Displaying Monitor Counters.”](#) You can clear ISL counters using the **perfMonitorClear** command, as described in [“Clearing Monitor Counters” on page 14-14.](#)

Monitoring Trunks

For trunked ISLs on Fabric OS v4.x switches, monitoring is set only on the master ISL, which communicates with the associated slave ISLs. For Fabric OS v3.x switches, monitoring can be set on slave ISLs.

End-to-end monitors are not supported for ISLs.

SilkWorm 3016, 3250, 3850, 3900, 12000, 24000, and 48000 models support eight filter-based monitors for trunks.

The SilkWorm 200E, 4012, and 4100 switches support 12 filter-based monitors for trunks.

Displaying Monitor Counters

Use the **perfMonitorShow** command to display the monitors on a specified port. For end-to-end counters, you can display either the cumulative count of the traffic detected by the monitors or a snapshot of the traffic at specified intervals.



Note

SilkWorm 200E, 4012, and 4100 outputs do not include CRC counts.

The command format is:

```
perfmonitorshow --class monitor_class [slotnumber/]portnumber [interval]
```

<i>monitor_class</i>	The monitor class, which can be one of EE (end-to-end), FLT (filter-based), or ISL (inter-switch link). The --class monitor_class operand is required.
<i>slotnumber</i>	Specifies the slot number for a SilkWorm 12000, 24000, or 48000 director. For all other switches, this operand is not required. The slot number must be followed by a slash (/) and the port number, so that each port is represented by both slot number (1 through 4 or 7 through 10) and port number (0 through 15). The SilkWorm director has a total of 10 slots. Slot numbers 5 and 6 are control processor blades; slots 1 through 4 and 7 through 10 are port blades. For 16-port blades, there are 16 ports, counted from the bottom, numbered 0 to 15. For 32-port blades, there are 32 ports numbered 0 to 31.
<i>portnumber</i>	Specifies a port number. Valid values for port number vary, depending on the switch type. This operand is required.
<i>interval</i>	Specifies an interval in seconds. The interval must be greater than or equal to five seconds. For end-to-end monitoring, the Tx and Rx counts are measured in bytes. This operand is optional.

Example: Displaying end-to-end monitor on a port at an interval of every 6 seconds

```
switch:admin> perfMonitorShow --class EE 4/5 6
perfmonitorshow 53, 6: Tx/Rx are # of bytes and crc is # of crc errors
      0          1          2          3          4
-----
crc   Tx   Rx  crc   Tx   Rx  crc   Tx   Rx  crc   Tx   Rx  crc   Tx   Rx
-----
0     0     0   0     0     0   0     0     0   0     0     0   0     0     0
0    53m   4.9m 0    53m   4.9m 0    53m   4.9m 0    53m   4.9m 0    53m   0
0    53m   4.4m 0    53m   4.4m 0    53m   4.4m 0    53m   4.4m 0    53m   0
0    53m   4.8m 0    53m   4.8m 0    53m   4.8m 0    53m   4.8m 0    53m   0
0    53m   4.6m 0    53m   4.6m 0    53m   4.6m 0    53m   4.6m 0    53m   0
0    53m   5.0m 0    53m   5.0m 0    53m   5.0m 0    53m   5.0m 0    53m   0
0    53m   4.8m 0    53m   4.8m 0    53m   4.8m 0    53m   4.8m 0    53m   0
0    53m   4.5m 0    53m   4.5m 0    53m   4.5m 0    53m   4.5m 0    53m   0
0    52m   4.5m 0    52m   4.5m 0    52m   4.5m 0    52m   4.5m 0    52m   0
0    52m   5.0m 0    52m   5.0m 0    52m   5.0m 0    52m   5.0m 0    52m   0
0    52m   4.5m 0    52m   4.5m 0    52m   4.5m 0    52m   4.5m 0    52m   0
0    52m   4.6m 0    52m   4.6m 0    52m   4.6m 0    52m   4.6m 0    52m   0
```

Example: Displaying EE monitors on a port

```
switch:admin> perfMonitorShow --class EE 4/5
There are 7 end-to-end monitor(s) defined on port 53.
```

KEY	SID	DID	OWNER_APP	OWNER_IP_ADDR	TX_COUNT	RX_COUNT	CRC_COUNT
0	0x58e0f	0x1182ef	TELNET	N/A	0x0000000000000000	0x0000000000000000	0x0000000000000000
0	0x21300	0x21dda	TELNET	N/A	0x00000004d0ba9915	0x0000000067229e65	0x0000000000000000
1	0x21300	0x21ddc	TELNET	N/A	0x00000004d0baa754	0x0000000067229e65	0x0000000000000000
2	0x21300	0x21de0	TELNET	N/A	0x00000004d0bab3a5	0x0000000067229e87	0x0000000000000000
3	0x21300	0x21de1	TELNET	N/A	0x00000004d0bac1e4	0x0000000067229e87	0x0000000000000000
4	0x21300	0x21de2	TELNET	N/A	0x00000004d0bad086	0x0000000067229e87	0x0000000000000000
5	0x11000	0x21fd6	WEB_TOOLS	192.168.169.40	0x00000004d0bade54	0x0000000067229e87	0x0000000000000000
6	0x11000	0x21fe0	WEB_TOOLS	192.168.169.40	0x00000004d0baed41	0x0000000067229e98	0x0000000000000000

To display filter-based monitor on a port at an interval of every 6 seconds

```
switch:admin> perfMonitorShow --class FLT 2/5 6
perfmonitorshow 21, 6
```

0	1	2	3	4	5	6
#Frames	#Frames	#Frames	#Frames	#Frames	#Frames	#Frames
0	0	0	0	0	0	0
26k	187	681	682	682	494	187
26k	177	711	710	710	534	176
26k	184	734	734	734	550	184
26k	182	649	649	649	467	182
26k	188	754	755	755	567	184
26k	183	716	716	717	534	183
26k	167	657	656	655	488	167
26k	179	749	749	749	570	179
26k	164	752	752	752	588	164
26k	190	700	700	700	510	190
26k	181	701	701	701	520	181
26k	200	750	750	751	550	201
26k	180	692	692	691	512	179
26k	179	696	696	696	517	179
26k	187	720	720	720	533	187
26k	200	722	722	722	522	200
26k	204	717	717	717	513	204

To display filter monitor information on a port

```
switch:admin> perfMonitorShow --class FLT 2/5
There are 7 filter-based monitors defined on port 21.
```

KEY	ALIAS	OWNER_APP	OWNER_IP_ADDR	FRAME_COUNT
0	SCSI_Frame	TELNET	N/A	0x000000000002c2229
1	SCSI_WR	TELNET	N/A	0x000000000000464a
2	SCSI_RW	TELNET	N/A	0x000000000000fd8c
3	SCSI_RW	WEB_TOOLS	192.168.169.40	0x00000000000007ba3
4	SCSI_RW	WEB_TOOLS	192.168.169.190	0x0000000000004f0e
5	SCSI_RD	WEB_TOOLS	192.168.169.40	0x0000000000002208
6	SCSI_WR	WEB_TOOLS	192.168.169.40	0x000000000000033a

To display ISL monitor information on a port

```
switch:admin> perfMonitorShow --class ISL 1/1
Total transmit count for this ISL: 1462326
Number of destination domains monitored: 3
Number of ports in this ISL: 2
Domain 97: 110379 Domain 98: 13965
Domain 99: 1337982
```

Clearing Monitor Counters

Before you clear statistics counters, verify the valid monitor numbers on a specific port using the **perfMonitorShow** command, to make sure the correct monitor counters are cleared. To clear statistics counters for all or a specified monitor, use the **perfMonitorClear** command. After the command has been executed, the telnet shell confirms that the counters on the monitor have been cleared.

The command format is:

```
perfmonitorclear --class monitor_class [slotnumber/]portnumber [monitorId]
```

<i>monitor_class</i>	The monitor class, which can be one of EE (end-to-end), FLT (filter-based), or ISL (inter-switch link). The --class <i>monitor_class</i> operand is required.
<i>slotnumber</i>	Specifies the slot number for a SilkWorm 12000, 24000, or 48000 director. For all other switches, this operand is not required. The slot number must be followed by a slash (/) and the port number, so that each port is represented by both slot number (1 through 4 or 7 through 10) and port number (0 through 15). The SilkWorm director has a total of 10 slots. Slot numbers 5 and 6 are control processor blades; slots 1 through 4 and 7 through 10 are port blades. For 16-port blades, there are 16 ports, counted from the bottom, numbered 0 to 15. For 32-port blades, there are 32 ports numbered 0 to 31.
<i>portnumber</i>	Specifies a port number. Valid values for port number vary, depending on the switch type. This operand is required.
<i>monitorId</i>	Specifies the monitor number to clear. Monitor numbers are defined when you create the monitor on a port. This operand is optional. If not specified, all monitor counters on the port are cleared. This operand does not apply to ISL monitors.



Note

In Fabric OS v3.1.0 and v4.1.0 (or later) the **portStatsClear** command clears AL_PA- based CRC error counters for all the ports in the same group.

To clear statistics counters for an end-to-end monitor

```
switch:admin> perfMonitorClear --class EE 1/2 5
End-to-End monitor number 5 counters are cleared

switch:admin> perfMonitorClear --class EE 1/2
This will clear ALL EE monitors' counters on port 2, continue?
(yes, y, no, n): [no] y
```

To clear statistics counters for a filter-based monitor

```
switch:admin> perfMonitorClear --class FLT 1/2 4
Filter-based monitor number 4 counters are cleared

switch:admin> perfMonitorClear --class FLT 1/2
This will clear ALL filter-based monitors' counters on port 2, continue? (yes, y, no,
y): [no] y
```

To clear statistics counters for an ISL monitor

```
switch:admin> perfMonitorClear --class ISL 1
This will clear ISL monitor on port 1, continue? (yes, y, no, n): [no] y
```

Saving and Restoring Monitor Configurations

To save the current end-to-end and filter monitor configuration settings into nonvolatile memory, use the **perfCfgSave** command:

```
switch:admin> perfCfgsave
This will overwrite previously saved Performance Monitoring settings in FLASH ROM.
Do you want to continue? (yes, y, no, n): [no] y
Please wait... Committing configuration...done.
Performance monitoring configuration saved in FLASH ROM.
```

To restore a saved monitor configuration, use the **perfCfgRestore** command. For example, to restore the original performance monitor configuration after making several changes:

```
switch:admin> perfCfgrestore
This will overwrite current Performance Monitoring settings in RAM. Do you want to
continue? (yes, y, no, n): [no] y
Please wait... Performance monitoring configuration restored from FLASH ROM.
```

To clear the previously saved performance monitoring configuration settings from nonvolatile memory, use the **perfCfgClear** command:

```
switch:admin> perfCfgclear
This will clear Performance Monitoring settings in FLASH ROM. The RAM settings won't
change. Do you want to continue? (yes, y, no, n): [no] y
Please wait... Committing configuration...done.
Performance Monitoring configuration cleared from FLASH.
```

Collecting Performance Data

Data collected through Advanced Performance Monitoring is deleted when the switch is rebooted. Using the Brocade Fabric Manager software application version 4.4.0 (or later), you can store performance data persistently. For details on this feature, refer to the *Fabric Manager Administrator's Guide*.

Configuring the PID Format

Port identifiers (called *PIDs*) are used by the routing and zoning services in Fibre Channel fabrics to identify ports in the network. All devices in a fabric must use the same PID format, so when you add new equipment to your SAN, you might need to change the PID format on legacy equipment.

About PIDs and PID Binding

The PID is a 24-bit address built from three 8-bit fields:

- domain
- area_ID
- AL_PA

Many scenarios cause a device to receive a new PID; for example, unplugging the device from one port and plugging it into a different port as part of fabric maintenance, or changing the domain ID of a switch, which might be necessary when merging fabrics, or changing compatibility mode settings.

Some device drivers use the PID to map logical disk drives to physical Fibre Channel counterparts. Most drivers can either change PID mappings dynamically (called *dynamic PID binding*) or use the WWN of the Fibre Channel disk for mapping (called *WWN binding*).

Some older device drivers behave as if a PID uniquely identifies a device (they use *static PID binding*). These device drivers should be updated, if possible, to use WWN or dynamic PID binding instead, because static PID binding creates problems in many routine maintenance scenarios. Fortunately, very few device drivers still behave this way. Many current device drivers enable you to select static PID binding as well as WWN binding. You should only select static binding if there is a compelling reason, and only after you have evaluated the impact of doing so.

Summary of PID Formats

SilkWorm switches employ these types of PID formats:

- VC encoded

This is the format defined by the SilkWorm 1000 series. Connections to these switches are not supported in Fabric OS v4.0.0 and later.

- native

Introduced with the SilkWorm 2000 series, this format supports up to 16 ports per switch.

- core

The default for SilkWorm 200E, 3016, 3250, 3850, 3900, 4012, and 4100 switches and SilkWorm 12000, 24000, and 48000 directors, this is the recommended format for Brocade switches and fabrics. It uses the entire 8-bit address space and directly uses the port number as the area_ID. It supports up to 256 ports per switch.

- extended edge

This format generates the same PID for a port on switches with 16 ports or less as would native PID format, but it also supports up to 256 ports per domain. It should be used only in cases where you cannot upgrade devices to dynamic PID binding and you absolutely cannot reboot your servers.

Extended edge PID is supported in Fabric OS v2.6.2 and later, v3.1.2 and later, and v4.2.0 and later.

In addition to the PID formats list here, Interoperability mode supports additional PID formats that are not discussed in this guide.

Impact of Changing the Fabric PID Format

If your fabric contains switches that use Native PID, it is recommended that you change the format to Core PID before you add the new, higher port count switches and directors. Also, it is recommended that you use Core PID when upgrading the Fabric OS version on 2000 and 3000 series switches.

Depending on your situation, the PID change might or might not entail fabric downtime:

- If you are running dual-fabrics with multipathing software, you can update one fabric at a time without disrupting traffic. Move all traffic onto one fabric in the SAN and update the other fabric. Then move the traffic onto the updated fabric, and update the final fabric.
- Without dual-fabrics, stopping traffic is highly recommended. This is the case for many routine maintenance situations, so dual-fabrics are always recommended for uptime-sensitive environments. If your fabric contains devices that employ static PID binding, or you do not have dual-fabrics, you must schedule downtime for the SAN to change the PID format.

You can find more details on the impact of PID changes in the following publications, which are available on the Brocade partner Web site. If you do not have access to this site, ask your support provider for these documents:

- *Brocade SilkWorm Design, Deployment, and Management Guide* (Publication Number: 53-0000366)
- *Brocade SAN Migration Guide* (Publication Number: 53-0000360)

The following sections describe various impacts of PID format changes in greater detail.

Host Reboots

In some Fibre Channel SAN environments, storage devices and host servers are bound to the host operating system by their PIDs (called their *Fibre Channel addresses*). In these environments, the hosts and target HBAs in a SAN need to know the full 24-bit PIDs of the hosts and targets they are communicating with, but they do not care how the PIDs are determined. But, if a storage device PID is changed, the host must reestablish a new binding, which requires the host to be rebooted.

With the introduction of the SilkWorm 200E, 3016, 3250, 3850, 3900, 4012, and 4100 switches and the SilkWorm 12000, 24000, and 48000 directors, the Native PID format used in earlier switches was supplemented with the Core PID format, which is capable of addressing higher port counts. Changing from Native PID format to Core PID format changes the PID, which requires hosts that use port binding to be rebooted.

Static PID Mapping Errors

If you can avoid using drivers that employ static PID binding, you should do so.

With the WWN or dynamic PID binding most typically used with drivers, changing the device's PID does not affect the PID mapping. However, before updating the PID format, it is necessary to determine whether any devices in the SAN use static PID binding.

For those few drivers that do use static PID binding, changing the PID format breaks the mapping, which must be fixed either by rebooting the host or by using a manual update procedure on the host.

To correct mapping errors caused by static PID binding, refer to the following sections:

- See [“Evaluating the Fabric” on page A-5](#) for details on finding devices that use static PID binding. Then refer to [“Online Update” on page A-8](#) or [“Offline Update” on page A-8](#) for recommendations.
- Refer to [“Converting Port Number to Area ID” on page A-12](#) for instructions.

Changes to Configuration Data

[Table A-1](#) lists various combinations of before-and-after PID formats, and indicates whether the configuration is affected.



Caution

After changing the fabric PID format, if the change invalidates the configuration data (see [Table A-1](#) to determine this), do not download old (pre-PID format change) configuration files to any switch on the fabric.

Table A-1 Effects of PID Format Changes on Configurations

PID Format Before Change	PID Format After Change	Configuration Effect?
Native	Extended Edge	No impact
Extended Edge	Native	No impact

Table A-1 Effects of PID Format Changes on Configurations (Continued)

PID Format Before Change	PID Format After Change	Configuration Effect?
Native	Core	You must: <ul style="list-style-type: none"> • Reenable zoning, if there is an active zone set and it uses port zones. • You do not need to reconfigure Destination ID (DID) if using: <ul style="list-style-type: none"> - Performance monitoring - The configure command to change the PID format - The performance monitor database - The zoning database - Security dcc database (if secure mode is enabled) The DID will be converted by FOS automatically.
Core	Native	
Extended Edge	Core	
Core	Extended Edge	

After changing the fabric PID format and verifying correct fabric operation, resave configuration data by running the **configUpload** command.

Before downgrading firmware, change the PID back to supported PIDs such as Core PID. If the database is automatically converted, save the converted database, and then download the older OS.

Selecting a PID format

All switches in a fabric must use the same PID format, so if you add a switch that uses a different PID format to a fabric, the switch will segment from the fabric. The format you select for your fabric depends on the mix of switches in the fabric, and to an extent on the specific releases of Fabric OS in use (for example, Extended Edge PID format is only available in Fabric OS v2.6.2 and later, Fabric OS v3.1.2 and later, and Fabric OS v4.2.0 and later).

If you are building a new fabric with switches running various Fabric OS versions, use Core PID format to simplify port-to-area_ID mapping.

[Table A-2](#) shows various combinations of existing fabrics, new switches added to those fabrics, and the recommended PID format for that combination. The criteria for the recommendations are first to

eliminate host reboots, and second to minimize the need for a host reboot in the future.

Table A-2 PID Format Recommendations For Adding New Switches

Existing Fabric OS Versions; PID Format	Switch to be Added	Recommendations (in Order of Preference)
v2.6.2 and later/v3.1.2 and later; Native PID	v2.6.2 and later/v3.1.2 and later	<ol style="list-style-type: none"> 1. Use Native PID format for new switch Host reboot is not required. 2. Convert existing fabric to Core PID format, upgrading the version of Fabric OS, if necessary. Set Core PID format for new switch. Host reboot <i>is</i> required. 3. If devices are bound statically and it is not possible to reboot, convert existing fabric to Extended Edge PID format, upgrading the version of Fabric OS, if necessary. Use Extended Edge PID format for new switch Host reboot is not required.
	v4.2.0 and later	<ol style="list-style-type: none"> 1. Convert existing fabric to Core PID format, upgrading the version of Fabric OS, if necessary. Set Core PID format for new switch. Host reboot <i>is</i> required. 2. If devices are bound statically and it is not possible to reboot, convert existing fabric to Extended Edge PID format, upgrading the version of Fabric OS, if necessary. Use Extended Edge PID format for new switch Host reboot is not required.
v2.6.2 and later/v3.1.2 and later/v4.2.0 and later; Core PID	v2.6.2 and later/v3.1.2 and later/v4.2.0 and later	<p>Use Core PID for new switch</p> <p>Host reboot is not required.</p>
v2.6.2 and later/v3.1.2 and later/v4.2.0 and later; Extended Edge PID	v2.6.2 and later/v3.1.2 and later/v4.2.0 and later	<p>Use Extended Edge PID for new switch</p> <p>Host reboot is not required.</p>

Evaluating the Fabric

In addition to this section, refer to the *Brocade SilkWorm Design, Deployment, and Management Guide* for information on evaluating the fabric.

If there is the possibility that your fabric contains host devices with static PID bindings, you should evaluate the fabric to:

- Find any devices that bind to PIDs
- Determine how each device driver will respond to the PID format change

- Determine how any multipathing software will respond to a fabric service interruption

If current details about the SAN are already available, it might be possible to skip the Data Collection step. If not, it is necessary to collect information about each device in the SAN. Any type of device might be able to bind by PID; each device should be evaluated before attempting an online update. This information has broad applicability, because PID-bound devices are not able to seamlessly perform in many routine maintenance or failure scenarios.

1. Collect device, software, hardware, and configuration data.

The following is a non-comprehensive list of information to collect:

- HBA driver versions
- Fabric OS versions
- RAID array microcode versions
- SCSI bridge code versions
- JBOD drive firmware versions
- Multipathing software versions
- HBA time-out values
- Multipathing software timeout values
- Kernel timeout values
- Configuration of switch

2. Make a list of manually configurable PID drivers.

Some device drivers do not automatically bind by PID, but allow the operator to manually create a PID binding. For example, persistent binding of PIDs to logical drives might be done in many HBA drivers. Make a list of all devices that are configured this way. If manual PID binding is in use, consider changing to WWN binding.

The following are some of the device types that might be manually configured to bind by PID:

- HBA drivers (persistent binding)
- RAID arrays (LUN access control)
- SCSI bridges (LUN mapping)

3. Analyze data.

After you have determined the code versions of each device on the fabric, they must be evaluated to find out if any automatically bind by PID. It might be easiest to work with the support providers of these devices to get this information. If this is not possible, you might need to perform empirical testing.

Binding by PID can create management difficulties in a number of scenarios. It is recommended that you not use drivers that bind by PID. If the current drivers do bind by PID, upgrade to WWN-binding drivers if possible.

The drivers shipping by default with HP/UX and AIX at the time of this writing still bind by PID, and so detailed procedures are provided for these operating systems in this chapter. Similar procedures can be developed for other operating systems that run HBA drivers that bind by PID.

There is no inherent PID binding problem with either AIX or HP/UX. It is the HBA drivers shipping with these operating systems that bind by PID. Both operating systems are expected to release HBA drivers that bind by WWN, and these drivers might already be available through some support channels. Work with the appropriate support provider to find out about driver availability.

It is also important to understand how multipathing software reacts when one of the two fabrics is taken offline. If the time-outs are set correctly, the failover between fabrics should be transparent to the users.

You should use the multipathing software to manually fail a path before starting maintenance on that fabric.

4. Perform empirical testing.

Empirical testing might be required for some devices, to determine whether they bind by PID. If you are not sure about a device, work with the support provider to create a test environment.

Create as close a match as practical between the test environment and the production environment, and perform an update using the procedure in [“Online Update” on page A-8](#).

Devices that bind by PID are unable to adapt to the new format, and one of three approaches must be taken with them:

- A plan can be created for working around the device driver’s limitations in such a way as to allow an online update. See the Detailed Procedures section for examples of how this could be done.
- The device can be upgraded to drivers that do not bind by PID.
- Downtime can be scheduled to reset the device during the core PID update process, which generally allows the mapping to be rebuilt.

If either of the first two options are used, the procedures should again be validated in the test environment.

Determine the behavior of multipathing software, including but not limited to:

- HBA time-out values
- Multipathing software time-out values
- Kernel time-out values

Planning the Update Procedure

Whether it is best to perform an offline or online update depends on the uptime requirements of the site.

- An offline update that all devices attached to the fabric be offline.
- With careful planning, it should be safe to update the core PID format parameter in a live, production environment. This requires dual fabrics with multipathing software. Avoid running backups during the update process, as tape drives tend to be very sensitive to I/O interruption. The online update process is only intended for use only in uptime-critical dual-fabric environments, with multipathing software (high-uptime environments should always use a redundant fabric SAN architecture). Schedule a time for the update when the least critical traffic is running.

All switches running any version of Fabric OS 3.1.2 and later or 4.2.0 and later are shipped with the Core Switch PID Format enabled, so it is not necessary to perform the PID format change on these switches.

Migrating from manual PID binding (such as persistent binding on an HBA) to manual WWN binding and upgrading drivers to versions that do not bind by PID can often be done before setting the core PID format. This reduces the number of variables in the update process.

Online Update

The following steps are intended to provide SAN administrators a starting point for creating site-specific procedures.

1. Back up all data and verify backups.
2. Verify that the multipathing software can automatically switchover between fabrics seamlessly. If there is doubt, use the software's administrative tools to manually disassociate or mark offline all storage devices on the first fabric to be updated.
3. Verify that I/O continues over the other fabric.
4. Disable all switches in the fabric to be updated, one switch at a time, and verify that I/O continues over the other fabric after each switch disable.
5. Change the PID format on each switch in the fabric.
6. Reenable the switches in the updated fabric one at a time. In a core/edge network, enable the core switches first.
7. After the fabric has reconverged, use the **cfgEnable** command to update zoning.
8. Update their bindings for any devices manually bound by PID. This might involve changing them to the new PIDs, or preferably changing to WWN binding.

For any devices automatically bound by PID, two options exist:

- a. Execute a custom procedure to rebuild its device tree online. Examples are provided in the [“Converting Port Number to Area ID” on page A-12](#) section of this chapter.
 - b. Reboot the device to rebuild the device tree. Some operating systems require a special command to do this, for example “boot -r” in Solaris.
9. For devices that do not bind by PID or have had their PID binding updated, mark online or reassociate the disk devices with the multipathing software and resume I/O over the updated fabric.
 10. Repeat with the other fabric(s).

Offline Update

The following steps are intended to provide SAN administrators a starting point for creating site-specific procedures.

1. Schedule an outage for all devices attached to the fabric.
2. Back up all data and verify backups.
3. Shut down all hosts and storage devices attached to the fabric.
4. Disable all switches in the fabric.
5. Change the PID format on each switch in the fabric.
6. Reenable the switches in the updated fabric one at a time. In a core/edge network, enable the core switches first.
7. After the fabric has reconverged, use the **cfgEnable** command to update zoning.

8. Bring the devices online in the order appropriate to the SAN. This usually involves starting up the storage arrays first, and the hosts last.
9. For any devices manually bound by PID, bring the device back online, but do not start applications. Update their bindings and reboot again if necessary. This might involve changing them to the new PIDs, or might (preferably) involve changing to WWN binding.
10. For any devices automatically bound by PID, reboot the device to rebuild the device tree (some operating systems require a special command to do this, such as “boot -r” in Solaris).
11. For devices that do not bind by PID or have had their PID binding updated, bring them back up and resume I/O.
12. Verify that all I/O has resumed correctly.

Hybrid Update

It is possible to combine the online and offline methods for fabrics where only a few devices bind by PID. Because any hybrid procedure is extremely customized, it is necessary to work closely with the SAN service provider in these cases.

Changing to Core PID Format

In Fabric OS release v4.2.0 and later, Native PID format is not supported; the default format is the Core PID format.

In Fabric OS v3.1.2 and later, Core PID format is the default configuration.

In Fabric OS v2.6.2 and later, Native PID format is the default configuration.

Although the PID format is listed in the configuration file, do not edit the file to change the setting there. Instead, use the CLI **configure** command. When you use the **configure** command, switch databases that contain PID-sensitive information are automatically updated. If you change the setting in the configuration file and then download the edited file, the PID format will be changed, but the database entries will not, and so they will be incorrect.

The following maps the PID format names to the names used in the management interfaces.

PID Format Name	Management Interface Name
native PID	switch PID address mode 0
core PID	switch PID address mode 1
extended edge PID	switch PID address mode 2

Before changing the PID format, determine if host reboots will be necessary. The section [“Host Reboots” on page A-2](#) summarizes the situations that might require a reboot.

Example

```

switch:admin> switchdisable
switch:admin> configure
Configure...

Fabric parameters (yes, y, no, n): [no] y

Domain: (1..239) [1]
BB credit: (1..27) [16]
R_A_TOV: (4000..120000) [10000]
E_D_TOV: (1000..5000) [2000]
WAN_TOV: (1000..120000) [0]
Data field size: (256..2112) [2112]
Sequence Level Switching: (0..1) [0]
Disable Device Probing: (0..1) [0]
Suppress Class F Traffic: (0..1) [0]
SYNC IO mode: (0..1) [0]
Switch PID Address Mode: (0..2) [1] < Set mode number here.
Per-frame Route Priority: (0..1) [0]
Long Distance Fabric: (0..1) [0]

```

Changing to Extended Edge PID Format

In rare cases, you might be affected by the presence in the fabric of drivers that rely on static binding to the dynamically assigned PID; for example, you might be installing a switch running Fabric OS v4.2.0 into a fabric consisting solely of Fabric OS v2.6.2 and later/v3.1.2 and later switches. In these cases, if you absolutely cannot reboot the affected servers when you upgrade your switches, you can choose Extended Edge PID format. It uses the same PID mapping for the first 16 ports and can support switches and directors with higher port counts. However, because Extended Edge format only supports 128 ports per domain, its use can lead to port addressing issues in directors.

Use the following procedure only if your fabric contains devices that are bound statically and you cannot reboot the host.

1. Determine if the current switch firmware versions meet the minimum supported version levels.

[Table A-3](#) lists the earliest Fabric OS version levels that support Extended Edge PID format. Use this table to determine if you need to upgrade the firmware in the switches in your fabric before you change the PID format.

Table A-3 Earliest Fabric OS Versions for Extended Edge PID format

Series 2000	SilkWorm 3200 and 3800 Switches	SilkWorm 3016, 3250, 3850, 3900, 12000, 24000 Switches	SilkWorm 4012 Switch	SilkWorm 200E, 48000 Switches
2.6.2	3.1.2	4.2.0	5.0.0	5.0.1

2. Update switch firmware as necessary.
 - a. Use the **fabricShow** command to verify the total number of switches in the fabric.
 - b. Download the correct firmware version to each switch as necessary.
 - c. Reboot all switches.

- d. Verify that the switches form a single fabric and that all domain IDs do not change after forming the fabric.
- e. Verify that the number of switches is the same.
3. Disable the switch by entering the **switchDisable** command.
4. Change the switch configuration in the fabric to Extended Edge PID format.
 - a. Configure Extended Edge PID (Format 2) on each switch. (See [Figure A-1](#) for a sample configure command on a SilkWorm switch running Fabric OS v3.1.2 and later and see [Figure A-1](#) for a sample configure command on a SilkWorm switch running Fabric OS 4.2.0 and later.)
 - b. Run the **switchEnable** command all switches.
 - c. Verify that all the switches form a fabric.
 - d. Use the **switchShow** command to verify the interswitch links (ISLs) are correct and the device links are correct.
 - e. Use the **fabricShow** command to verify that the number of switches are the same as those when starting this procedure.
 - f. Use the **nsAllShow** command to verify the total number of devices is the same as those when starting this procedure.
 - g. For dual fabrics, repeat steps 1 through 5 for the other fabric.

Example: Configure Command on a Switch Running Fabric OS 3.1.2

```
Configure...

Fabric parameters (yes, y, no, n): [no] yes

  Domain: (1..239) [217]
  BB credit: (1..27) [16]
  R_A_TOV: (4000..120000) [10000]
  E_D_TOV: (1000..5000) [2000]
  Data field size: (256..2112) [2112]
  Sequence Level Switching: (0..1) [0]
  Disable Device Probing: (0..1) [0]
  Suppress Class F Traffic: (0..1) [0]
  SYNC IO mode: (0..10) [0]
  Switch PID Format : (0..2) [0] 2
  Per-frame Route Priority: (0..1) [0]
  Long Distance Fabric: (0..1) [0]

Virtual Channel parameters (yes, y, no, n): [no] ^D
Committing configuration...done.
0x102fd500 (tshell): Apr 15 16:53:31
  WARNING CONFIG-PIDCHANGE_DISPLACE, 3, Switch PID format changed to Extended Edge
PID Format
```

Example: Configure Command on a Switch Running Fabric OS 5.0.1

```

configure

Configure...

Fabric parameters (yes, y, no, n): [no] y

Domain: (1..239) [11]
R_A_TOV: (4000..120000) [10000]
E_D_TOV: (1000..5000) [2000]
WAN_TOV: (0..30000) [0]
MAX_HOPS: (7..19) [7]
Data field size: (256..2112) [2112]
Sequence Level Switching: (0..1) [0]
Disable Device Probing: (0..1) [0]
Suppress Class F Traffic: (0..1) [0]
Switch PID Format: (1..2) [1] 2
Per-frame Route Priority: (0..1) [0]
Long Distance Fabric: (0..1) [0]
BB credit: (1..27) [16]

Insistent Domain ID Mode (yes, y, no, n): [no]
Virtual Channel parameters (yes, y, no, n): [no]
Zoning Operation parameters (yes, y, no, n): [no]
RSCN Transmission Mode (yes, y, no, n): [no]
Arbitrated Loop parameters (yes, y, no, n): [no]
System services (yes, y, no, n): [no]
Portlog events enable (yes, y, no, n): [no]
ssl attributes (yes, y, no, n): [no]
http attributes (yes, y, no, n): [no]
snmp attributes (yes, y, no, n): [no]
rpcd attributes (yes, y, no, n): [no]
cfgload attributes (yes, y, no, n): [no]
webtools attributes (yes, y, no, n): [no]

Switch PID format has changed to PID Format 2 ('Extended Edge PID')

```

Converting Port Number to Area ID

Except for the following cases, the area ID is equal to the port number:

- when you perform a port swap operation
- when you enable Extended Edge (also known as “displaced PID”) PID on the director

If you are using Extended Edge PID format (for example, the SilkWorm 48000 with configuration option 5) and would like to map the output of the port number to the area ID, use the following formula (for ports 0-127):

$$a = (p + 16) \% 128$$

where:

a area

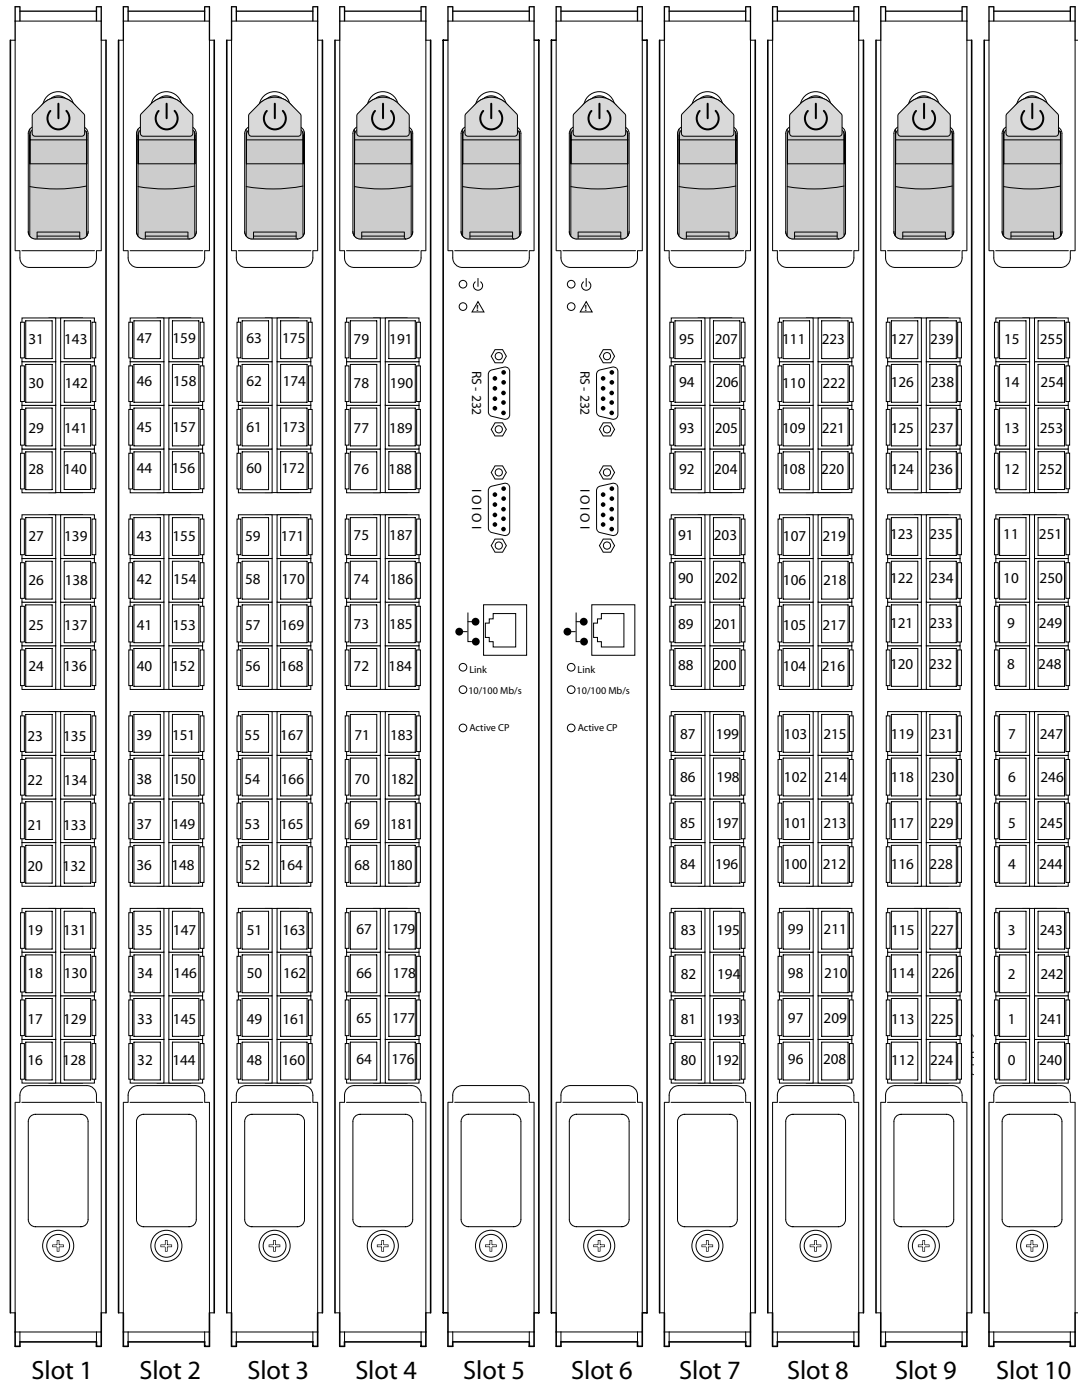
p port number

$\%$ modulus (or remainder)

$$0 \leq p < 128$$

When the port number is greater than or equal to 128, the area ID and port number are the same.

[Figure A-1](#) shows a SilkWorm 48000 with Extended Edge PID.

Table A-4 SilkWorm 48000 With Extended Edge PID

Performing PID Format Changes

There are several routine maintenance procedures which might result in a device receiving a new PID. Examples include, but are not limited to:

- Changing compatibility mode settings
- Changing switch domain IDs
- Merging fabrics
- Relocating devices to new ports or new switches (that is, for Add, Move, Change type operations)
- Updating the core PID format
- Using hot spare switch ports to deal with failures

In every case where devices employ static PID binding, any such procedure becomes difficult or impossible to execute without downtime.

In some cases, device drivers allow you to specify static PID binding. In these cases, such devices must be identified and their PID binding should be changed to WWN binding.

The following sections contain a basic procedure that summarizes the steps necessary to perform PID format changes without disrupting the fabric, and special procedures for HP/UX and AIX.

Basic Procedure

This process should be executed as part of the overall online or offline update process. However, it can be implemented in a stand-alone manner on a non-production fabric, or a switch that has not yet joined a fabric.

1. Ensure that all switches in the fabric are running Fabric OS versions that support the addressing mode. It is recommended that you use v2.6.2 for SilkWorm 2000 series switches, v3.1.2 for SilkWorm 3200 and 3800 switches, v4.2.0 for SilkWorm 12000 and 24000 directors, as well as SilkWorm 3016, 3250, 3850, and 3900 switches, and v5.0.1 for SilkWorm 48000.



Note

All switches running any version of Fabric OS 4.0.0 and later are shipped with the Core Switch PID Format enabled, so it is not necessary to perform the PID format change on these switches.

2. Telnet into one of the switches in the fabric.
3. Disable the switch by entering the **switchDisable** command.
4. Enter the **configure** command (the configure prompts display sequentially).
5. Enter **y** after the “Fabric parameters” prompt.
6. Enter **1** at the “Core Switch PID Format” prompt.
7. Complete the remaining prompts or press **Ctrl-d** to accept the remaining settings without completing all the prompts.
8. Repeat steps 2 through 7 for the remaining switches in the fabric.
9. Reenable the switch by entering the **switchEnable** command.

Example

```
switch:admin> switchdisable
switch:admin> configure
Configure...
Fabric parameters (yes, y, no, n): [no] yes
Domain: (1..239) [1]
R_A_TOV: (4000..120000) [10000]
E_D_TOV: (1000..5000) [2000]
Data field size: (256..2112) [2112]
Sequence Level Switching: (0..1) [0]
Disable Device Probing: (0..1) [0]
Suppress Class F Traffic: (0..1) [0]
SYNC IO mode: (0..1) [0]
Core Switch PID Format: (0..2) [0] 1
Per-frame Route Priority: (0..1) [0]
Long Distance Fabric: (0..1) [0]
BB credit: (1..27) [16]
```

10. After all switches are updated to use the new PID format and reenabled, verify that the fabric has fully reconverged (each switch “sees” the other switches).
11. Enter **cfgEnable [active_zoning_config]** on one of the switches in the fabric to update zoning to use the new PID form.
This does not change the definition of zones in the fabric, but merely causes the lowest level tables in the zoning database to be updated with the new PID format setting. It is only necessary to do this once per fabric; the zoning update automatically propagates to all switches.
At this point, all switches in the fabric are operating in the new addressing mode.

HP/UX Procedure

This procedure is not intended to be comprehensive. It provides a starting point from which a SAN administrator could develop a site-specific procedure for a device that binds automatically by PID, and cannot be rebooted due to uptime requirements.

1. Backup all data. Verify backups.
2. If you are not using multipathing software, stop all I/O going to all volumes connected through the switch/fabric to be updated.
3. If you are not using multipathing software, unmount the volumes from their mount points using **umount**. The proper usage would be **umount <mount_point>**. For example:

```
umount /mnt/jbod
```
4. If you are using multipathing software, use that software to remove one fabric’s devices from its configuration.
5. Deactivate the appropriate volume groups using **vgchange**. The proper usage would be **vgchange -a n <path_to_volume_group>**. For example:

```
vgchange -a n /dev/jbod
```
6. Make a backup copy of the volume group directory using **tar** from within /dev. For example:

```
tar -cf /tmp/jbod.tar jbod
```

7. Export the volume group using `vgexport`. The proper usage would be **`vgexport -m <mapfile> <path_to_volume_group>`**. For example:

```
vgexport -m /tmp/jbod_map /dev/jbod
```

8. Connect to each switch in the fabric
9. Issue the **`switchDisable`** command.
10. Issue the **`configure`** command and change the Core Switch PID Format to 1.
11. Issue the command **`cfgEnable [effective_zone_configuration]`**. For example:

```
cfgEnable my_zones
```

12. Issue the **`switchEnable`** command. Enable the core switches first, then the edges.
13. Clean the `lvmtab` file by using the command `vgscan`.
14. Change to `/dev` and untar the file that was tared in step 4. For example:

```
tar -xf /tmp/jbod.tar
```

15. Import the volume groups using `vgimport`. The proper usage would be **`vgimport -m <mapfile> <path_to_volume_group> <physical_volume_path>`**. For example:

```
vgimport -m /tmp/jbod_map /dev/jbod /dev/dsk/c64t8d0 /dev/dsk/c64t9d0
```

16. Activate the volume groups using `vgchange`. The proper usage would be **`vgchange -a y <path_to_volume_group>`**. For example:

```
vgexport -a y /dev/jbod
```

17. If you are not using multipathing software, mount all devices again and restart I/O. For example:
- ```
mount /mnt/jbod
```
18. If you are using multipathing software, reenable the affected path. The preceding steps do not “clean up” the results from `ioscan`. When viewing the output of `ioscan`, notice the that the original entry is still there, but now has a status of `NO_HW`.

| # ioscan -funC disk |     |                         |                  |                    |                    |             |
|---------------------|-----|-------------------------|------------------|--------------------|--------------------|-------------|
| Class               | I   | H/W Path                | Driver S/W State | H/W Type           | Description        |             |
| -----               |     |                         |                  |                    |                    |             |
| disk                | 0   | 0/0/1/1.2.0             | adisk CLAIMED    | DEVICE             | SEAGATE ST39204LC  |             |
|                     |     |                         | /dev/dsk/clt2d0  | /dev/rdisk/clt2d0  |                    |             |
| disk                | 1   | 0/0/2/1.2.0             | adisk CLAIMED    | DEVICE             | HP                 | DVD-ROM 304 |
|                     |     |                         | /dev/dsk/c3t2d0  | /dev/rdisk/c3t2d0  |                    |             |
| disk                | 319 | 0/4/0/0.1.2.255.14.8.0  | adisk CLAIMED    | DEVICE             | SEAGATE ST336605FC |             |
|                     |     |                         | /dev/dsk/c64t8d0 | /dev/rdisk/c64t8d0 |                    |             |
| disk                | 320 | 0/4/0/0.1.18.255.14.8.0 | adisk NO_HW      | DEVICE             | SEAGATE ST336605FC |             |
|                     |     |                         | /dev/dsk/c65t8d0 | /dev/rdisk/c65t8d0 |                    |             |

19. To remove the original (outdated) entry, the command `rmsf` (remove special file) will be needed. The proper usage for this command would be **`rmsf -a -v <path_to_device>`**. For example:

```
rmsf -a -v /dev/dsk/c65t8d0
```

20. Validate that the entry has been removed by using the command **ioscan -funC disk**. In this example, the NO\_HW entry is no longer listed:

```
het46 (HP-50001)> ioscan -funC disk
```

| Class | I   | H/W Path               | Driver S/W State | H/W Type           | Description        |
|-------|-----|------------------------|------------------|--------------------|--------------------|
| disk  | 0   | 0/0/1/1.2.0            | adisk CLAIMED    | DEVICE             | SEAGATE ST39204LC  |
|       |     |                        | /dev/dsk/clt2d0  | /dev/rdisk/clt2d0  |                    |
| disk  | 1   | 0/0/2/1.2.0            | adisk CLAIMED    | DEVICE             | HP DVD-ROM 304     |
|       |     |                        | /dev/dsk/c3t2d0  | /dev/rdisk/c3t2d0  |                    |
| disk  | 319 | 0/4/0/0.1.2.255.14.8.0 | adisk CLAIMED    | DEVICE             | SEAGATE ST336605FC |
|       |     |                        | /dev/dsk/c64t8d0 | /dev/rdisk/c64t8d0 |                    |

21. Repeat for all fabrics.
22. Issue the **switchEnable** command. Enable the core switches first, then the edges.

## AIX Procedure

This procedure is not intended to be comprehensive. It provides a starting point from which a SAN administrator can develop a site-specific procedure for a device that binds automatically by PID, and cannot be rebooted due to uptime requirements.

1. Backup all data. Verify backups.
2. If you are not using multipathing software, stop all I/O going to all volumes connected through the switch or fabric to be updated.
3. If you are not using multipathing software, varyoff the volume groups. The command usage is **varyoffvg <volume\_group\_name>**. For example:

```
varyoffvg datavg
```

4. If you are not using multipathing software, unmount the volumes from their mount points using **umount**. The command usage is **umount <mount\_point>**. For example:

```
umount /mnt/jbod
```

5. If you are using multipathing software, use that software to remove one fabric's devices from its configuration.
6. Remove the device entries for the fabric you are migrating. For example, if the HBA for that fabric is fcs0, execute the command:

```
rmdev -Rdl fcs0
```

7. Connect to each switch in the fabric.
8. Issue the **switchDisable** command.
9. Issue the **configure** command and change the Core Switch PID Format to 1.
10. Issue the **configEnable [effective\_zone\_configuration]** command. For example:

```
configenable my_config
```

11. Issue the **switchEnable** command. Enable the core switches first, then the edges.
12. Rebuild the device entries for the affected fabric using the **cfgMgr** command. For example:

```
cfgmgr -v
```

This command might take several minutes to complete.

13. If you are not using multipathing software, vary on the disk volume groups. The proper usage would be

**varyonvg <volume\_group\_name>**. For example:

```
varyonvg datavg
```

14. If you are not using multipathing software, mount all devices again and restart I/O. For example:

```
mount /mnt/jbod
```

15. If you are using multipathing software, reenable the affected path.
16. Repeat for all fabrics.

## Swapping Port Area IDs

If a device that uses port binding is connected to a port that fails, you can use port swapping to make another physical port use the same PID as the failed port. The device can then be plugged into the new port without the need to reboot the device.

Use the following procedure to swap the port area IDs of two physical switch ports. In order to swap port area IDs, the port swap feature must be enabled, and both switch ports must be disabled. The swapped area IDs for the two ports remain persistent across reboots, power cycles, and failovers.

Swap area IDs for a pair of switch ports as follows:

1. Connect to the switch and log in as admin.
2. Enable the port swap feature:

```
portswapenable
```

3. **SilkWorm 200E, 3016, 3250, 3850, 3900, 4012, and 4100 switches:** Enter the following commands:

```
portdisable port1
portdisable port2
```

**SilkWorm 12000, 24000, and 48000 directors:** Enter the following commands:

```
portdisable slot/port1
portdisable slot/port2
```

4. **SilkWorm 200E, 3016, 3250, 3850, 3900, 4012, and 4100 switches:** Enter the following command:

```
portswap port1 port2
```

**SilkWorm 12000, 24000, and 48000 directors:** Enter the following command:

```
portswap slot1/port1 slot2/port2
```

5. Verify that the port area IDs have been swapped:

```
portswapshow
```

A table is shown that shows the physical port numbers and the logical area IDs for any swapped ports.

6. Disable the port swap feature:

```
portswapdisable
```



# Configuring Interoperability Mode

---

This topic provides information on setting up a heterogeneous fabric that includes Brocade switches and switches from other manufacturers.

The *interoperability* mode enables Brocade switches and others to exchange interoperability parameters, allowing their fabrics to merge into one fabric with one principal switch and unique domain IDs.

The **interopMode** command must be executed on all Brocade switches in the fabric. Each Brocade switch must be rebooted after changing interoperability mode.

Switches from other manufacturers might also require the execution of one or more commands to set up interoperability mode. Refer to their documentation for details.

In a heterogeneous fabric, each Brocade switch must have interoperability mode enabled (see [“Enabling and Disabling Interoperability Mode” on page B-5](#)).

To provide maximum compatibility between switches, several features are not available in heterogeneous fabrics (see [“Supported Brocade Features” on page B-2](#) and [“Unsupported Brocade Features” on page B-2](#)).

## Vendor Switch Requirements

Interoperability has been tested on McDATA switches.

McDATA ES-3016 or ED-5000 models, or an equivalent OEM version that is plug-compatible, is required. Contact your switch supplier for the supported switches and firmware versions.

To determine whether or not a mixed-vendor SAN is supported, work with the switch provider to determine if your SAN design is valid. Important variables that determine the supportability of a particular mixed vendor SAN include the number of switches, version of Fabric OS, the topology, number of ISLs, number of connected devices, and hop count.

## Brocade Switch Requirements

Following are the SilkWorm software requirements:

- 2000-series models must be running Fabric OS v2.6.0 or later.
- 3000-series models must be running Fabric OS v3.1.0 or later.
- 3016, 3250, 3850, 3900, 12000, and 24000 models must be running 4.2.0 or later.
- 4100 model must be running Fabric OS v4.4.0 or later.
- 4012 model must be running Fabric OS v5.0.0 or later.
- 200E model running Fabric OS v5.0.1 or later.
- 48000 model must be running Fabric OS v5.0.1 or later.
- A zoning license and a fabric license must be installed on each Brocade switch.

## Supported Brocade Features

The following features are supported on Brocade switches in interoperability mode:

- Brocade Fabric Watch
- Brocade Fabric Access API functions

Accessible from Brocade switches only, but switch information for non-Brocade switches is reported. The object information and zoning actions are configurable from the API.

- Brocade translative mode

Registers private storage target devices into the fabric, it can be used in a heterogeneous fabric if the devices are connected directly to Brocade switches. The devices will be accessible from any port on the fabric.



---

**Note**

Switches with a condor ASIC do not support translative mode.

---

## Unsupported Brocade Features

In a heterogeneous fabric, the following Brocade optional features are not supported and cannot be installed on any Brocade switch in the fabric:

- Extended Edge PID format
- Quickloop and QuickLoop Zoning
- Secure Fabric OS
- Timer Server function
- Open E\_Port
- Broadcast Zoning
- Management Server Service and FDMI

- QuickLoop Fabric Assist
- Remote Switch
- Extended Fabrics
- Trunking
- Alias Server
- Platform Service
- Virtual Channels
- FC-IP

## Configuration Recommendations

The following is recommended when configuring an interoperable fabric:

- Avoid domain ID conflicts before fabric reconfiguration. Every switch in the fabric must have a unique domain ID.
- When you are configuring multiple switches, you should wait for a fabric reconfiguration after adding or removing each switch.

## Configuration Restrictions

In interoperable fabrics, the following restrictions apply:

- Do not use Extended Edge PID mode.
- There is an architecture maximum of 31 switches.
- Domain IDs must be in the 97 to 127 value range for successful connection to McDATA switches. The firmware automatically assigns a valid domain ID, if necessary, when the **interopmode** command is enabled on the switch.
- The **fabricShow** command only shows the WWN and domain ID for McDATA switches. No IP address or switch name information is provided. Brocade switches show all parameters.
- If you are managing zoning from Brocade switches, then all Brocade switches must have at least one direct connection to another Brocade switch. For example, you cannot have a McDATA switch between two Brocade switches if you are managing zoning from the Brocade switches.
- LC IBM GBICs are not supported if they are connected to a McDATA ISL.
- When a Brocade switch gets a new domain ID assigned through a fabric reconfiguration, the new domain ID is written to nonvolatile memory and the old domain ID value is overwritten. When a McDATA switch gets a new domain ID assigned through a fabric reconfiguration, it keeps the original domain ID in nonvolatile memory.

In this scenario, when the domain ID of both a McDATA switch and a Brocade switch are changed via fabric reconfiguration, on the next and subsequent fabric reconfiguration(s), the Brocade switch attempts to use the new ID (from the nonvolatile memory) while McDATA attempts to use its old ID (from the nonvolatile memory). This situation might cause a domain ID overlap to occur during multiple fabric reconfigurations. Domain ID overlap is not supported for Brocade/McDATA interoperability.

- Between Brocade switches, you can connect more than one ISL when in interoperability mode.

## Zoning Restrictions

The following restrictions apply to zoning in interoperable fabrics:

- When interoperability mode is in effect, the space available for the zoning database is about half the usual size. The maximum zoning database size in interoperability mode is 1:1.8 of the native mode zoning database size
- Only zoning by port WWN is allowed; you must use the port WWN of the device, such as 10:00:00:00:c9:28:c7:c6.
- Zone members specified by node WWN are ignored.
- Zone configurations that use either physical port numbers or port IDs are not supported in interoperability mode. Zoning using port numbers uses the actual physical port numbers on the switch; for example slot 1, port 5.
- When a zoning configuration is not in effect, by default all ports are isolated and traffic is not permitted. This is unlike Brocade behavior where *Interoperability* mode is off (and all data traffic is enabled).
- SilkWorm 200E, 3016, 3200, 3250, 3800, 3850, 3900, 4012, 4100, 12000, 24000, and 48000 models provide hardware enforcement of the port WWN zones only for devices attached to their ports. Devices attached to end-ports on non-Brocade switches or Brocade 2000-series switches are enforced by Name Server (soft) zoning only.
- Web Tools can be used for zone configuration as long as Brocade switches are connected directly to each other. If Web Tools is used to set up zoning, then it must be used as the only zone management method.
- Brocade switches connected to a McDATA switch receive the effective configuration when a zone merge occurs. (McDATA only has an effective zone configuration and discards the defined zone configuration when it sends merge information to the Brocade switch.) However, a zone update sends both the defined and the effective configuration to all switches.
- When a SilkWorm switch or director is reconfiguring, wait until the fabric routes are completely set up before entering zoning commands that must propagate to other switches. Use the **fabricShow** command to verify that all fabric routes are set up and all switch IP addresses and names are present. (The **fabricShow** command only shows the WWN and domain ID for switches from other manufacturers.)
- The maximum number of items that can be stored in the zoning configuration database depends on the switches in the fabric, whether or not interoperability mode is enabled, and the number of bytes required for each item.

You can use the **cfgSize** command to check both the maximum available size and the currently saved size. If you believe you are approaching the maximum, you can save a partially completed zoning configuration and use the **cfgSize** command to determine the remaining space

## Zone Name Restrictions

The name field must contain the ASCII characters that actually specify the name, not including any required fill bytes. Names must follow these rules:

- Length must be between 1 and 64 characters.
- All characters must be 7-bit ASCII.
- The first character must be a letter, which can be either upper case (A-Z) or lower case (a-z).
- Any character other than the first character must be lower case (a-z), upper case (A-Z), a single-digit number (0-9), dash (-), or underscore (\_).

## Enabling and Disabling Interoperability Mode

Before enabling interoperability mode, inspect the individual fabrics for compatibility.

- Make sure that zones meet the zone criteria and restrictions listed in [“Zoning Restrictions” on page B-4](#)).
- Use the **msplMgmtDeactivate** command to disable any platform management functions. (Refer to the *Fabric OS Command Reference Manual* for additional command information).

### To enable interoperability mode

1. Verify that you have implemented all the Brocade prerequisites necessary to enable interoperability mode on the fabric (refer to [“Configuration Recommendations”](#) and [“Configuration Restrictions” on page B-3](#)).
2. Connect to the switch and log in as admin.
3. Enter the **switchDisable** command to disable the switch.
4. Use the **configure** command to set the domain ID to a number in the range from 97 to 126. (For detailed instructions, refer to [“Working With Domain IDs” on page 2-15](#).)
5. Enter the **interopmode 1** command to enable interoperability. This command resets a number of parameters and enables interactive mode.
6. Reboot the switch after changing the interoperability mode:

```
switch:admin> switchdisable
switch:admin> configure
Configure...
Fabric Parameters (yes, y, no, n): [no] y
Domain (1...239): [1] 97
...

switch:admin> interopmode 1
The switch effective configuration will be lost when the operating mode is changed;
do you want to continue? (yes, y, no, n): [no] y

done.
Interopmode is enabled

Note: It is recommended that you reboot this switch for the new change to take
effect.
switch:admin>
```

7. Repeat [step 2](#) through [step 6](#) on all Brocade switches in the fabric.
8. Enable interoperability on each non-Brocade switch. (Refer to the switch documentation.)

9. After enabling interoperability mode on all switches, physically connect the non-Brocade switches into the Brocade fabric, one at a time.

### To disable interoperability mode

1. Connect to the switch and log in as admin.
2. Enter the **switchDisable** command to disable the switch.
3. Enter the **interopmode 0** command to disable interoperability. This command resets a number of parameters and disables interactive mode.
4. Reboot the switch after changing the interoperability mode:

```
switch:admin> switchdisable
switch:admin> interopmode 0
The switch effective configuration will be lost when the operating mode is changed;
do you want to continue? (yes, y, no, n): [no] y

done.
Interopmode is disabled

Note: It is recommended that you reboot this switch for the new change to take effect.
switch:admin>
```

5. Wait for a fabric reconfiguration after removing each switch.
6. Each non-Brocade switch might require the execution of a similar command to disable interoperability.
7. Repeat this procedure on all Brocade switches in the fabric.

# Using Remote Switch

---

The Brocade Remote Switch feature, which aids in ensuring gateway compatibility, was formerly a licensed feature. Its functionality is now available as part of the Fabric OS standard feature set through the use of the **portCfgIslMode** command, which is described in [“Linking Through a Gateway” on page 2-17](#). For those who use Remote Switch as part of their legacy set of tools, this appendix contains a description and procedure for the feature.

Remote Switch, enables you to connect two remote Brocade switches over an IP network, enabling communication of IP or ATM protocols as well as Fibre Channel traffic.

The Brocade Remote Switch feature functions with the aid of a “bridging device” or Fibre Channel gateway. The gateway supports both a Fibre Channel physical interface and a secondary, non-Fibre Channel physical interface, such as IP, SONET, or ATM. Remote Switch functions over E\_Port connections. With Remote Switch on both fabrics, the gateway accepts Fibre Channel frames from one fabric, tunnels them across the network, and passes them to the other fabric. From the viewpoint of the connected hosts and storage devices, fabrics using Remote Switch interact the same as locally connected switches.

Remote Switch provides many of the same capabilities of normal ISL links including,

- Coordinated fabric services

The Remote Switch fabric configuration fully supports all fabric services, including distributed name service, registered state change notification, and alias service.

- Distributed management

Management tools such as Advanced Web Tools, Fabric OS, and SNMP are available from both the local switch and the remote switch. Switch management is routed through the Fibre Channel connection; thus, no additional network connection is required between sites.

- Support for interswitch links (ISLs)

Sites requiring redundant configurations can connect multiple E\_Ports to remote sites by using multiple gateways. Standard Fabric OS routing facilities automatically maximize throughput and provide automatic failover during interruption on the WAN connection.

The Remote Switch feature operates in conjunction with a gateway. The gateway provides an E\_Port interface that links to the SilkWorm E\_Port. After the link between the two E\_Ports has been negotiated, the gateway E\_Port moves to passthrough mode and passes Fibre Channel traffic from the SilkWorm E\_Port to the WAN.

The gateway accepts Fibre Channel frames from one side of a Remote Switch fabric, transfers them across a WAN, and passes them to the other side of the Remote Switch fabric.

Remote Switch can be used for these types of gateway devices:

- Fibre Channel over ATM
- Fibre Channel over IP
- Fibre Channel over SONET
- Fibre Channel over DWDM

Most of these gateway devices have enough buffers to cover data transfer over a wide area network (WAN). The SilkWorm switches on each side of the gateway must have identical configurations. Only qualified SFPs should be used.

You must connect the fabrics through the gateway device, and make sure that the **configure** parameters are compatible with the gateway device.

You might be required to reconfigure the following parameters, depending on the gateway requirements:



#### Note

Consult your gateway vendor for supported and qualified configurations.

- R\_A\_TOV: Specify a Resource Allocation Timeout Value compatible with your gateway device.
- E\_D\_TOV: Specify a Error Detect Timeout Value compatible with your gateway device
- Data field size: Specify the maximum Fibre Channel data field reported by the fabric. Verify the maximum data field size the network-bridge can handle. Some bridges might not be able to handle a maximum data field size of 2112.
- BB credit: Specify the number of Buffer-to-Buffer credits for Nx\_port devices.
- Suppress Class F Traffic: Use this parameter to disable class F traffic. Some network-bridge devices might not have a provision for handling class F frames. In this case, the transmission of class F frames must be suppressed throughout the entire Remote Switch fabric.

To set the access and reconfigure these parameters:

1. Connect to the switch and log in as admin.
2. Enter the **switchDisable** command to disable the switch.
3. Enter the **configure** command.
4. At the Fabric Parameters prompt enter **yes**.
5. Press **Enter** to scroll through the **Fabric Parameters** without changing their values, until you reach the parameter you want to modify.
6. Specify a new parameter value that is compatible with your gateway device.
7. Press **Enter** to scroll through the remainder of the configuration parameters. Make sure that the configuration changes are committed to the switch.
8. Repeat for all switches in the fabrics to be connected through a gateway device. These parameters must be identical on each switch in the fabric, and between fabrics connected through the gateway device.

### Example

This example shows how to modify the data field size and suppress class F traffic on a switch:

```
switch:admin> switchdisable
switch:admin> configure
Configure...
Fabric parameters (yes, y, no, n): [no] yes
Domain: (1..239) [3]
 R_A_TOV: (4000..120000) [10000]
 E_D_TOV: (1000..5000) [2000]
 Data field size: (256..2112) [2112] 1000
 Sequence Level Switching: (0..1) [0]
 Disable Device Probing: (0..1) [0]
 Suppress Class F Traffic: (0..1) [0] 1
 VC Encoded Address Mode: (0..1) [0]
 Per-frame Route Priority: (0..1) [0]
 Long Distance Fabric: (0..1) [0]
 BB credit: (1..16) [16]
Virtual Channel parameters (yes, y, no, n): [no]
 Zoning Operation parameters (yes, y, no, n): [no]
 RSCN Transmission Mode (yes, y, no, n): [no]
 NS Operation Parameters (yes, y, no, n): [no]
 Arbitrated Loop parameters (yes, y, no, n): [no]
 System services (yes, y, no, n): [no]
 Portlog events enable (yes, y, no, n): [no]
Committing configuration...done.
switch:admin>
```



# Understanding Legacy Password Behavior

The following sections provide password information for early versions of Fabric OS firmware.

## Password Management Information

The following table describes the password standards and behaviors between various versions of firmware.

**Table D-1** Account/Password Characteristics Matrix

| Topic                                                                                                                                                                                             | v4.0.0                     | v4.1.0 to v4.2.0                                                                           | v4.4.0 to 5.0.1                                                                                  |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------|--------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|
| Number of default accounts on the switch                                                                                                                                                          | 4, chassis based           | <b>12000</b> - 8 for the chassis, 4 per switch<br><b>3016, 3250, 3850, 3900, 24000</b> - 4 | <b>12000</b> - 8 for the chassis, 4 per switch<br><b>3016, 3250, 3850, 3900, 4100, 24000</b> - 4 |
| Default account names                                                                                                                                                                             | root, factory, admin, user | root, factory, admin, user. 3016: administrative account is called USERID.                 | root, factory, admin, user. 3016: administrative account is called USERID.                       |
| Account name changing feature                                                                                                                                                                     | No                         | No, regardless of security mode.                                                           | <b>3016</b> - supports the <b>userRename</b> command.                                            |
| Maximum and minimum amount of characters for a password                                                                                                                                           | 0 - 8 (Standard UNIX)      | 8 - 40 characters with printable ASCII                                                     | 8 - 40 characters with printable ASCII                                                           |
| Can different switch instances use a different password for the same account login level? For example, the password for admin for switch 0 can be different from password for admin for switch 1. | No                         | Yes for 12000 switch.<br>n/a for all other switches.                                       | Yes for 12000 switch.<br>n/a for all other switches.                                             |
| Does the root account use restricted shell?                                                                                                                                                       | No                         | No                                                                                         | No                                                                                               |

**Table D-1** Account/Password Characteristics Matrix (Continued)

| Topic                                                                                               | v4.0.0                                                                                                                                              | v4.1.0 to v4.2.0                                                                                                                                                                                                                                                                                                                  | v4.4.0 to 5.0.1                                                                                                                                                                                                                                                                                                                   |
|-----------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| When connecting to a factory installed switch, do you use the default passwords?                    | Yes                                                                                                                                                 | Yes                                                                                                                                                                                                                                                                                                                               | Yes                                                                                                                                                                                                                                                                                                                               |
| Does a user need to know the old passwords when changing passwords using the <b>passwd</b> command? | Yes, except when the root user changes another user's password. This is standard UNIX behavior; Fabric OS does not enforce any additional security. | Old password is required only when changing password for the same level user password. Changing password for lower level user does not require old password. For example, users connect as admin; old admin password is required to change the admin password. But old user password is not required to change the user password. | Old password is required only when changing password for the same level user password. Changing password for lower level user does not require old password. For example, users connect as admin; old admin password is required to change the admin password. But old user password is not required to change the user password. |
| Can <b>passwd</b> change higher-level passwords? For example, can admin change root password?       | Yes, but will ask for the "old password" of the higher-level account (example "root").                                                              | Yes; if users connect as admin, they can change the root, factory, and admin passwords. However, if one connects as user, one can only change the user password.                                                                                                                                                                  | Yes, if users connect as admin, they can change the root, factory and admin passwords after first entering the old password for the respective account.                                                                                                                                                                           |
| Can API change passwords?                                                                           | Yes, only for admin.                                                                                                                                | Yes, only for admin.                                                                                                                                                                                                                                                                                                              | Yes, only for admin.                                                                                                                                                                                                                                                                                                              |
| Can Web Tools change passwords?                                                                     | No                                                                                                                                                  | No                                                                                                                                                                                                                                                                                                                                | No                                                                                                                                                                                                                                                                                                                                |
| Can SNMP change passwords?                                                                          | No                                                                                                                                                  | No                                                                                                                                                                                                                                                                                                                                | No                                                                                                                                                                                                                                                                                                                                |

# Password Prompting Behaviors

The following table describes the expected password prompting behaviors of various Fabric OS versions.

**Table D-2** Password Prompting Matrix

| Topic                                                                                                               | v4.0.0                                                                                                                                                                      | v4.1.0 and later                                                                                                                                                            |
|---------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Must <i>all</i> password prompts be completed for <i>any</i> change to take effect?                                 | No. Partial changes of all four passwords are allowed.                                                                                                                      | No. Partial changes of all four passwords are allowed.                                                                                                                      |
| When does the password prompt appear?                                                                               | When users connect as root, factory, or admin, the accounts with default password will be prompted for change. The accounts with non-default password will NOT be prompted. | When users connect as root, factory, or admin, the accounts with default password will be prompted for change. The accounts with non-default password will NOT be prompted. |
| Is a user forced to answer password prompts before getting access to the firmware?                                  | No, users can type in <b>Ctrl-c</b> to get out of password prompting.                                                                                                       | No, users can type in <b>Ctrl-c</b> to get out of password prompting.                                                                                                       |
| Do users need to know the old root password when answering prompting?                                               | Yes in v4.0.0<br>*No in v4.0.2 only                                                                                                                                         | No                                                                                                                                                                          |
| Are new passwords forced to be set to something different than the old passwords?                                   | Yes                                                                                                                                                                         | Yes                                                                                                                                                                         |
| Is password prompting disabled when security mode is enabled?                                                       | Yes                                                                                                                                                                         | Yes                                                                                                                                                                         |
| Is the <b>passwd</b> command disabled until the user has answered password prompting?                               | No                                                                                                                                                                          | Yes                                                                                                                                                                         |
| Does password prompting reappear when passwords are changed back to default using the <b>passwd</b> command?        | Yes                                                                                                                                                                         | No                                                                                                                                                                          |
| Does password prompting reappear when passwords are changed back to default using the <b>passwdDefault</b> command? | Yes                                                                                                                                                                         | Yes                                                                                                                                                                         |

# Password Migration During Firmware Changes

Table D-3 describes the expected outcome of password settings when upgrading or downgrading firmware for various Fabric OS versions.

**Table D-3** Password Migration Behavior During Firmware Upgrade/Downgrade

| Topic                                                                                    | v4.2.0 to 5.0.1                                                                                                                                                                                                       | v4.4.0 to v5.0.1                                                                                                                                                           |
|------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Passwords used when upgrading to a newer firmware release for the first time.            | Default accounts and passwords are preserved.                                                                                                                                                                         | Default accounts and passwords are preserved.                                                                                                                              |
| Passwords preserved during subsequent firmware upgrades                                  | MUA accounts created during a previous upgrade to 4.2.0 are restored.                                                                                                                                                 | MUA accounts created during a previous upgrade to 5.0 are restored.                                                                                                        |
| Passwords used if downgrading to an older firmware for the first time                    | Downgrades to 4.2.0 preserve existing default accounts and restore any disabled default accounts. The default account passwords existing immediately prior to the downgrade are preserved. MUA accounts are disabled. | Downgrades to 4.4.0 preserve all existing default accounts, MUA accounts and passwords. MUA accounts with the switchAdmin role have the same permissions as the user role. |
| When downgrading to an older firmware at subsequent times, which passwords will be used? | Downgrades to 4.2.0 preserve existing default accounts and restore any disabled default accounts. The default account passwords existing immediately prior to the downgrade are preserved. MUA accounts are disabled. | Downgrades to 4.4.0 preserve all existing default accounts, MUA accounts and passwords. MUA accounts with the switchAdmin role have the same permissions as the user role. |
| When downgrading then upgrading again, what passwords will be used?                      | Any MUA accounts created during the original upgrade are restored. All passwords are unchanged.                                                                                                                       | All default and MUA accounts and passwords remain unchanged.                                                                                                               |

# Password Recovery Options

The following table describes the options available when one or more types of passwords are lost.

**Table D-4** Password Recovery Options

| Topic                                                                                                                                 | v4.0.0                                                                         | v4.1.0 and later                                                                                                                                                                                                     |
|---------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| If all the passwords are forgotten, what is the password recovery mechanism? Are these procedures non-disruptive recovery procedures? | Contact your switch service provider. A non-disruptive procedure is available. | Contact your switch service provider. A non-disruptive procedure is available.                                                                                                                                       |
| If a user has only the root password, what is the password recovery mechanism?                                                        | Root can change any password by using the <b>passwd</b> command.               | Use <b>passwd</b> command to set other passwords.<br>Use <b>passwdDefault</b> command to set all passwords to default.                                                                                               |
| How to recover boot PROM password?                                                                                                    | n/a                                                                            | Contact your switch service provider and provide the recovery string.<br><br>Refer to <a href="#">“Setting the Boot PROM Password” on page 3-34</a> for instructions on setting the password with a recovery string. |
| How do I recover a user, admin, or factory password?                                                                                  | Refer to <a href="#">“Recovering Forgotten Passwords” on page 3-39</a> .       |                                                                                                                                                                                                                      |



# Zone Merging Scenarios

Table E-1 provides information on merging zones and the expected results.

**Table E-1** Zone Merging Scenarios

| Description                                                                                                                            | Switch A                                                  | Switch B                                               | Expected Results                                                                                                                      |
|----------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------|--------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| <b>Switch A</b> with a defined configuration<br><b>Switch B</b> does not have a defined configuration                                  | defined:<br>cfg1:<br>zone1: ali1; ali2<br>effective: none | defined: none<br>effective: none                       | Configuration from <b>Switch A</b> to propagate throughout the fabric in an inactive state, because the configuration is not enabled. |
| <b>Switch A</b> with a defined and enabled configuration<br><b>Switch B</b> has a defined configuration but no effective configuration | defined: cfg1<br>zone1: ali1; ali2<br>effective: cfg1:    | defined: cfg1<br>zone1: ali1; ali2<br>effective: none  | Configuration from <b>Switch A</b> to propagate throughout the fabric. The configuration is enabled after the merge in the fabric.    |
| <b>Switch A</b> and <b>Switch B</b> have the same defined configuration. Neither have an enabled configuration.                        | defined: cfg1<br>zone1: ali1; ali2<br>effective: none     | defined: cfg1<br>zone1: ali1; ali2<br>effective: none  | No change (clean merge).                                                                                                              |
| <b>Switch A</b> and <b>Switch B</b> have the same defined and enabled configuration.                                                   | defined: cfg1<br>zone1: ali1; ali2<br>effective: cfg1:    | defined: cfg1<br>zone1: ali1; ali2<br>effective: cfg1: | No change (clean merge).                                                                                                              |
| <b>Switch A</b> does not have a defined configuration<br><b>Switch B</b> with a defined configuration                                  | defined: none<br>effective: none                          | defined:cfg1<br>zone1: ali1; ali2<br>effective: none   | <b>Switch A</b> will absorb the configuration from the fabric.                                                                        |
| <b>Switch A</b> does not have a defined configuration<br><b>Switch B</b> with a defined configuration                                  | defined: none<br>effective: none                          | defined:cfg1<br>zone1: ali1; ali2<br>effective: cfg1   | <b>Switch A</b> will absorb the configuration from the fabric - with cfg1 as the effective cfg.                                       |
| <b>Switch A</b> and <b>Switch B</b> have the same defined configuration. Only <b>Switch B</b> has an enabled configuration.            | defined: cfg1<br>zone1: ali1; ali2<br>effective: none     | defined: cfg1<br>zone1: ali1; ali2<br>effective: cfg1  | Clean merge - with cfg1 as the effective cfg.                                                                                         |

**Table E-1** Zone Merging Scenarios (Continued)

| Description                                                                                                              | Switch A                                                                   | Switch B                                                                   | Expected Results                                                                                                                                   |
|--------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------|----------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Switch A</b> and <b>Switch B</b> have different defined configurations. Neither have an enabled zone configuration.   | defined: cfg2<br>zone2: ali3; ali4<br>effective: none                      | defined: cfg1<br>zone1: ali1; ali2<br>effective: none                      | Clean merge - the new cfg will be a composite of the two --<br>defined: cfg1<br>zone1: ali1; ali2<br>cfg2:<br>zone2: ali3; ali4<br>effective: none |
| <b>Switch A</b> and <b>Switch B</b> have different defined configurations. <b>Switch B</b> has an enabled configuration. | defined: cfg2<br>zone2: ali3; ali4<br>effective: none                      | defined: cfg1<br>zone1: ali1; ali2<br>effective: cfg1                      | Clean merge - The new cfg will be a composite of the two, with cfg1 as the effective cfg.                                                          |
| Effective cfg mismatch                                                                                                   | defined: cfg1<br>zone1: ali1; ali2<br>effective: cfg1<br>zone1: ali1; ali2 | defined: cfg2<br>zone2: ali3; ali4<br>effective: cfg2<br>zone2: ali3; ali4 | Fabric segments due to: Zone Conflict cfg mismatch                                                                                                 |
| cfg content mismatch                                                                                                     | defined: cfg1<br>zone1: ali1; ali2<br>effective: irrelevant                | defined: cfg1<br>zone1: ali3; ali4<br>effective: irrelevant                | Fabric segments due to: Zone Conflict content mismatch                                                                                             |
|                                                                                                                          | defined: cfg1<br>zone1: ali1; ali2<br>effective: irrelevant                | defined: cfg1<br>zone1: ali1; ali4<br>effective: irrelevant                | Fabric segments due to: Zone Conflict content mismatch                                                                                             |
| Same content - different effective cfg name                                                                              | defined: cfg1<br>zone1: ali1; ali2<br>effective: cfg1<br>zone1: ali1; ali2 | defined:cfg2<br>zone1: ali1; ali2<br>effective:cfg2<br>zone1: ali1; ali2   | Fabric segments due to: Zone Conflict cfg mismatch                                                                                                 |
| Same content - different zone name                                                                                       | defined: cfg1<br>zone1: ali1; ali2<br>effective: irrelevant                | defined: cfg1<br>zone2: ali1; ali2<br>effective: irrelevant                | Fabric segments due to: Zone Conflict content mismatch                                                                                             |
| Same content - different alias name                                                                                      | defined: cfg1 ali1: A; B<br>effective: irrelevant                          | defined:cfg1:ali2: A; B<br>effective: irrelevant                           | Fabric segments due to: Zone Conflict content mismatch                                                                                             |
| Same name - different types                                                                                              | effective: zone1:<br>MARKETING                                             | effective: cfg1:<br>MARKETING                                              | Fabric segments due to: Zone Conflict type mismatch                                                                                                |
| Same name - different types                                                                                              | effective: zone1:<br>MARKETING                                             | effective: alias1:<br>MARKETING                                            | Fabric segments due to: Zone Conflict type mismatch                                                                                                |
| Same name - different types                                                                                              | effective:cfg1:<br>MARKETING                                               | effective:alias1:<br>MARKETING                                             | Fabric segments due to: Zone Conflict type mismatch                                                                                                |

# Upgrading Firmware in Single CP Mode

---

For all SilkWorm switches and directors, the **firmwareDownload** command, by default, performs a full installation, automatic reboot (*autoreboot*), and automatic firmware commit (*autocommit*). Automatic reboot and automatic commit modes are not selectable by default; however, they become selectable when *single CP mode* is enabled by entering the *-s* option on the command line. In this case, **firmwareDownload** disables autoreboot and continues to enable *autocommit* mode by default.



---

**Note**

You should only use the following procedures if instructed to do so by your service provider.

---

Your service provider might ask you to perform this procedure on the on your SilkWorm switch—or on one or both CPs in director models—under these circumstances:

- To prevent the firmware commit that occurs after downloading, so that you can restore previous versions.
- To control the timing of the execution of the *haReboot* command, so that you can prestage the firmware ahead of time.

Always ensure that your switch is in a steady state (for example, no fabric configuration or ISL/cable changes) before entering the *haReboot* command. The *haReboot* command synchronizes firmware versions and associated data in real time. If your switch is not in a steady state when you enter the *haReboot* command, it will likely perform a cold restart for the affected area; this can disrupt the ASIC and machine traffic.

Your service provider may ask you to perform the single CP blade procedure on SilkWorm directors if a CP blade fails and the replacement CP blade is running a version of firmware that cannot synchronize with the current active CP blade.

For information about messages that might appear during the procedures, refer to the *Fabric OS System Error Message Reference Manual*.

## To upgrade SilkWorm 3016, 3250, 3850, 3900, and 4100 models:

Specify the *-s* option on the command line for single CP mode. You will be prompted for other options and can enable *autocommit* and *autoreboot* using the following procedure.

1. Connect to the switch and log in as admin.
2. Enter the following command:

```
firmwaredownload -s
```

Be sure to type a space between the command and the *-s* option.

3. Enter the IP address of the FTP server where the firmware is stored.

4. Enter your user name for the server.
5. Enter the full path to the firmware file on the server; for example:

```
/pub/v5.0.1/release.plist
```

6. Enter your password.
7. Answer the next prompts as indicated here:

```
Do Auto Commit after reboot [Y]: y
```

If you specify no, you must manually enter the **firmwareCommit** command.

```
Reboot system after download [N]: y
```

The default is no. If you take the default, you must later use the **haReboot** command to perform a high-availability reboot manually.

In v4.4.0 or later, the *Full Install* option is not available.

8. Wait for the firmware download to finish. (Start a new telnet session and use the **firmwareDownloadStatus** command to check the status.)

#### Example

```
switch: admin> firmwaredownload -s
Server Name or IP Address: 192.1.2.3
User Name: JohnDoe
File Name: /pub/v5.0.1/release.plist
Password: *****
Do Auto Commit after reboot [Y]: y
Reboot system after download [N]: y
Firmwaredownload has started.
.
.
.
```

### To upgrade a single Silkworm 12000/24000 CP blade

Although it is possible to upgrade firmware on one CP blade at a time, you should not do so under normal circumstances because it might disrupt director switch operations if not executed in the proper sequence, or if the version of firmware is down-level and part of an unsupported version.

When the two CP blades are not running the same firmware versions, it might be necessary to disable one or the other to maintain fabric stability. For information on the commands used to achieve this, refer to the **haDisable** and **haFailover** commands in the *Fabric OS Command Reference Manual*.

The following procedure allows you to upgrade a single CP blade. This procedure can be used with Fabric OS v4.0.0d and later.

1. Connect to the switch and log in as admin.
2. Enter the **haShow** command to determine which CP blade is the active and which one is the standby. In the following example, the active CP blade is CP0, and the standby CP blade is CP1:

```
switch:admin> hashow
Local CP (Slot 5, CP0): Active
Remote CP (Slot 6, CP1): Standby, Healthy
HA enabled, Heartbeat Up, HA State in sync
```

3. Log in to the standby CP blade as admin.
4. Enter the following command to upgrade a new version of the firmware to the standby CP blade:

```
firmwaredownload -s
```

Be sure to type a space between the command and the -s option.

5. Enter the IP address of the FTP server where the firmware is stored.
6. Enter your user name for the server.
7. Enter the full path to the firmware file on the server; for example:

```
/pub/v5.0.1/release.plist
```

8. Enter your password.
9. Answer the next prompts as indicated here:

```
Do Auto Commit after reboot [Y]: y
```

If you answer no, you must manually enter the **firmwareCommit** command.

```
Reboot system after download [N]: y
```

The default is no. If you take the default, you must later use the **haReboot** command to perform a high-availability reboot manually.

After you upgrade to v4.4.0 or later, the *Full Install* option is no longer available.

The default is no. If you take the default, you must manually reboot the CP.

### Example

```
switch: admin> firmwaredownload -s
Server Name or IP Address: 192.1.2.3
User Name: JohnDoe
File Name: /pub/v5.0.1/release.plist
Password: *****
Full Install (Otherwise upgrade only) [Y]: y
Do Auto Commit after reboot [Y]: y
Reboot system after download [N]: y
```

10. Wait for the firmware download to finish. (Start a new telnet session and use the **firmwareDownloadStatus** command to check the status.)
11. Enter the **haShow** command to verify that the two CP blades are synchronized.

12. Reboot the standby CP blade (if you set the option to reboot automatically to no in [step 9](#)).
13. Log in to the same CP blade and enter the **firmwareDownloadStatus** command to verify firmware has downloaded successfully and has either committed or is in the process of doing so. (If you set the option to do auto commit after reboot to no in [step 9](#), you must enter the **firmwareCommit** command manually.)
14. Enter the **haFailover** command to make the current CP blade active (with the updated firmware).
15. Repeat [step 1](#) through [step 13](#) on the other CP blade if the firmware versions are different.

# Glossary

---

## #

**8b/10b encoding** An encoding scheme that converts each 8-bit byte into 10 bits. Used to balance 1s and 0s in high-speed transports.

## A

**ABTS** Abort Basic Link Service. Also referred to as “Abort Sequence.”

**ACC** Accept link service reply. The normal reply to an Extended Link Service request (such as FLOGI), indicating that the request has been completed.

**address identifier** A 24-bit or 8-bit value used to identify the source or destination of a frame. Refer to S\_ID and DID.

**AL\_PA** Arbitrated loop physical address. A unique 8-bit value assigned during loop initialization to a port in an arbitrated loop. Alternately, “arbitrated loop parameters.”

**AL\_TIME** Arbitrated loop timeout value. Twice the amount of time it would take for a transmission word to propagate around a worst-case loop. The default value is 15 milliseconds (ms).

**alias** A logical grouping of elements in a fabric. An alias is a collection of port numbers and connected devices, used to simplify the entry of port numbers and WWNs when creating zones.

**alias address identifier** An address identifier recognized by a port in addition to its standard identifier. An alias address identifier can be shared by multiple ports. *See also* “[alias](#)”.

**alias AL\_PA** An AL\_PA value recognized by an L\_Port in addition to the AL\_PA assigned to the port. *See also* “[AL\\_PA](#)”.

**alias server** A fabric software facility that supports multicast group management.

**ARB** Arbitrative primitive signal. Applies only to an arbitrated loop topology. Transmitted as the fill word by an L\_Port to indicate that the port is arbitrating access to the loop.

**arbitrated loop** A shared 100-MB/sec Fibre Channel transport structured as a loop. Can support up to 126 devices and one fabric attachment. *See also* “[topology](#)”.

**arbitration** A method of gaining orderly access to a shared-loop topology.

|                       |                                                                                                                                                                                                                                                                                                                                    |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>area number</b>    | In Brocade Fabric OS v4.0.0 and above, ports on a switch are assigned a logical area number. Port area numbers can be viewed by entering the <b>switchshow</b> command. They are used to define the operative port for many Fabric OS commands: for example, area numbers can be used to define the ports within an alias or zone. |
| <b>ARR</b>            | Asynchronous response router. Refers to Management Server GS_Subtype Code E4, which appears in <b>portlogdump</b> command output.                                                                                                                                                                                                  |
| <b>ASD</b>            | Alias server daemon. Used for managing multicast groups by supporting the create, add, remove, and destroy functions.                                                                                                                                                                                                              |
| <b>ASIC</b>           | Application-specific integrated circuit. <i>[Necessary? Basic. –ed.]</i>                                                                                                                                                                                                                                                           |
| <b>ATM</b>            | Asynchronous Transfer Mode. A transport used for transmitting data over LANs or WANs that transmit fixed-length units of data. Provides any-to-any connectivity and allows nodes to transmit simultaneously.                                                                                                                       |
| <b>authentication</b> | The process of verifying that an entity in a fabric (such as a switch) is what it claims to be. <i>See also</i> “ <a href="#">digital certificate</a> ”, “ <a href="#">switch-to-switch authentication</a> ”.                                                                                                                      |
| <b>autocommit</b>     | A feature of the <b>firmwaredownload</b> command. Enabled by default, <b>autocommit</b> commits new firmware to both partitions of a control processor.                                                                                                                                                                            |
| <b>autoreboot</b>     | Refers to the <b>-b</b> option of the <b>firmwaredownload</b> command. Enabled by default.                                                                                                                                                                                                                                         |

## B

|                                    |                                                                                                                                                                                                                                                                 |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>BB_Credit</b>                   | Buffer-to-buffer credit. The number of frames that can be transmitted to a directly connected recipient or within an arbitrated loop. Determined by the number of receive buffers available. <i>See also</i> “ <a href="#">buffer-to-buffer flow control</a> ”. |
| <b>beacon</b>                      | A tool in which all of the port LEDs on a switch are set to flash from one side of the switch to the other, to enable identification of an individual switch in a large fabric. A switch can be set to beacon by a CLI command or through Brocade Web Tools.    |
| <b>beginning running disparity</b> | The disparity at the transmitter or receiver when the special character associated with an ordered set is encoded or decoded. <i>See also</i> “ <a href="#">disparity</a> ”.                                                                                    |
| <b>BIST</b>                        | Built-in self-test.                                                                                                                                                                                                                                             |
| <b>bit synchronization</b>         | The condition in which a receiver is delivering retimed serial data at the required bit error rate.                                                                                                                                                             |
| <b>block</b>                       | As it applies to Fibre Channel technology, upper-level application data that is transferred in a single sequence.                                                                                                                                               |
| <b>bloom</b>                       | The code name given to the third-generation Brocade Fabric ASIC.                                                                                                                                                                                                |

|                                      |                                                                                                                                                                             |
|--------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>broadcast</b>                     | The transmission of data from a single source to all devices in the fabric, regardless of zoning. <i>See also</i> <a href="#">“multicast”</a> , <a href="#">“unicast”</a> . |
| <b>buffer-to-buffer flow control</b> | Management of the frame transmission rate in either a point-to-point topology or in an arbitrated loop. <i>See also</i> <a href="#">“BB_Credit”</a> .                       |
| <b>bypass circuitry</b>              | Circuits that automatically remove a device from the data path when valid signals are dropped.                                                                              |

## C

|                         |                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>CA</b>               | Certificate authority. A trusted organization that issues digital certificates. <i>See also</i> <a href="#">“digital certificate”</a> .                                                                                                                                                                                                                                                                                                 |
| <b>CAM</b>              | Content-addressable memory.                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Class 1 service</b>  | The class of frame-switching service for a dedicated connection between two communicating ports (also called "connection-oriented service"). Includes acknowledgement of frame delivery or nondelivery.                                                                                                                                                                                                                                 |
| <b>Class 2 service</b>  | A connectionless class of frame-switching service that includes acknowledgement of frame delivery or nondelivery.                                                                                                                                                                                                                                                                                                                       |
| <b>Class 3 service</b>  | A connectionless class of frame-switching service that does not include acknowledgement of frame delivery or nondelivery. Can be used to provide a multicast connection between the frame originator and recipients, with acknowledgement of frame delivery or nondelivery.                                                                                                                                                             |
| <b>Class 4 service</b>  | A connection-oriented service that allows fractional parts of the bandwidth to be used in a virtual circuit.                                                                                                                                                                                                                                                                                                                            |
| <b>Class 6 service</b>  | A connection-oriented multicast service geared toward video broadcasts between a central server and clients.                                                                                                                                                                                                                                                                                                                            |
| <b>Class F service</b>  | The class of frame-switching service for a direct connection between two switches, allowing communication of control traffic between the E_Ports. Includes acknowledgement of data delivery or nondelivery.                                                                                                                                                                                                                             |
| <b>class of service</b> | A specified set of delivery characteristics and attributes for frame delivery.                                                                                                                                                                                                                                                                                                                                                          |
| <b>CLS</b>              | Close primitive signal. Used only in an arbitrated loop. Sent by an L_Port that is currently communicating in the loop, to close communication to another L_Port.                                                                                                                                                                                                                                                                       |
| <b>configuration</b>    | <p>(1) A set of parameters that can be modified to fine-tune the operation of a switch. Use the <b>configshow</b> command to view the current configuration of your switch.</p> <p>(2) In Brocade Zoning, a zoning element that contains a set of zones. The Configuration is the highest-level zoning element and is used to enable or disable a set of zones on the fabric. <i>See also</i> <a href="#">“zone configuration”</a>.</p> |
| <b>COS</b>              | Class of service.                                                                                                                                                                                                                                                                                                                                                                                                                       |

|               |                                                                                                                                                                    |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>CP</b>     | Control processor.                                                                                                                                                 |
| <b>credit</b> | As it applies to Fibre Channel technology, the number of receive buffers available to transmit frames between ports. <i>See also</i> <a href="#">“BB_Credit”</a> . |

## D

|                                   |                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>D_ID</b>                       | Destination identifier. A 3-byte field in the frame header, used to indicate the address identifier of the N_Port to which the frame is headed.                                                                                                                                                                       |
| <b>defined zone configuration</b> | The set of all zone objects defined in the fabric. Can include multiple zone configurations. <i>See also</i> <a href="#">“zone configuration”</a> .                                                                                                                                                                   |
| <b>digital certificate</b>        | An electronic document issued by a CA (certificate authority) to an entity, containing the public key and identity of the entity. Entities in a secure fabric are authenticated based on these certificates. <i>See also</i> <a href="#">“authentication”</a> , <a href="#">“CA”</a> , <a href="#">“public key”</a> . |
| <b>disparity</b>                  | The proportion of 1s and 0s in an encoded character. "Neutral disparity" means an equal number of each, "positive disparity" means a majority of 1s, and "negative disparity" means a majority of 0s.                                                                                                                 |
| <b>DLS</b>                        | Dynamic load-sharing. Dynamic distribution of traffic over available paths. Allows for recomputing of routes when an Fx_Port or E_Port changes status.                                                                                                                                                                |
| <b>domain controller</b>          | A domain controller (or embedded port) communicates with and gets updates from other switches' embedded ports. The well-known address is <i>fffdd</i> , where <i>dd</i> = domain number).                                                                                                                             |
| <b>domain ID</b>                  | A unique identifier for all switches in a fabric, used in routing frames. Usually automatically assigned by the principal switch but can be assigned manually. The domain ID for a Brocade SilkWorm switch can be any integer from 1 through 239.                                                                     |

## E

|                      |                                                                                                                                                                                                                                                                                                                |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>E_D_TOV</b>       | Error-detect timeout value. The minimum amount of time a target waits for a sequence to complete before initiating recovery. Can also be defined as the maximum time allowed for a round-trip transmission before an error is declared. <i>See also</i> <a href="#">“R_A_TOV”</a> , <a href="#">“RR_TOV”</a> . |
| <b>E_Port</b>        | Expansion port. A type of switch port that can be connected to an E_Port on another switch to create an ISL. <i>See also</i> <a href="#">“ISL”</a> .                                                                                                                                                           |
| <b>ELP</b>           | Exchange link parameters.                                                                                                                                                                                                                                                                                      |
| <b>ELS</b>           | Extended link service. ELSs are sent to the destination N_Port to perform the requested function or service. ELS is a Fibre Channel standard that is sometimes referred to as "Fibre Channel Physical (FC_PH) ELS."                                                                                            |
| <b>EM</b>            | Environmental monitor. Monitors FRUs and reports failures.                                                                                                                                                                                                                                                     |
| <b>embedded port</b> | An embedded port (or domain controller) communicates and get updates from other switches' embedded ports. The well-known address is <i>fffdd</i> , where <i>dd</i> = domain number.                                                                                                                            |

|                     |                                                                                                                                                                            |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>entry fabric</b> | The basic Brocade software license that allows one E_Port per switch.                                                                                                      |
| <b>EOF</b>          | End of frame. A group of ordered sets used to mark the end of a frame.                                                                                                     |
| <b>error</b>        | As it applies to the Fibre Channel industry, a missing or corrupted frame, timeout, loss of synchronization, or loss of signal (link errors).                              |
| <b>exchange</b>     | The highest-level Fibre Channel mechanism used for communication between N_Ports. Composed of one or more related sequences, it can work in either one or both directions. |

## F

|                        |                                                                                                                                                                                                                                        |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>F_BSY</b>           | Fabric port busy frame. A frame issued by the fabric to indicate that a frame cannot be delivered because the fabric or destination N_Port is busy.                                                                                    |
| <b>F_Port</b>          | Fabric port. A port that is able to transmit under fabric protocol and interface over links. Can be used to connect an N_Port to a switch. <i>See also</i> “ <a href="#">FL_Port</a> ”, “ <a href="#">Fx_Port</a> ”.                   |
| <b>F_RJT</b>           | Fabric port reject frame. A frame issued by the fabric to indicate that delivery of a frame is being denied, perhaps because a class is not supported, there is an invalid header, or no N_Port is available.                          |
| <b>fabric</b>          | A Fibre Channel network containing two or more switches in addition to hosts and devices. Also referred to as a “switched fabric.” <i>See also</i> “ <a href="#">SAN</a> ”, “ <a href="#">topology</a> ”.                              |
| <b>Fabric Manager</b>  | An optionally licensed Brocade software. Fabric Manager is a GUI that allows for fabric-wide administration and management. Switches can be treated as groups, and actions such as firmware downloads can be performed simultaneously. |
| <b>fabric name</b>     | The unique identifier assigned to a fabric and communicated during login and port discovery.                                                                                                                                           |
| <b>fabric services</b> | Codes that describe the communication to and from any well-known address.                                                                                                                                                              |
| <b>fabric topology</b> | The arrangement of switches that form a fabric.                                                                                                                                                                                        |
| <b>Fabric Watch</b>    | An optionally licensed Brocade software. Fabric Watch can be accessed through either the command line or Advanced Web Tools, and it provides the ability to set thresholds for monitoring fabric conditions.                           |
| <b>failover</b>        | Describes the Brocade SilkWorm 12000 process of one CP blade passing active status to another CP blade. A failover is nondisruptive.                                                                                                   |
| <b>FAN</b>             | Fabric address notification. Retains the AL_PA and fabric address when a loop reinitializes, if the switch supports FAN.                                                                                                               |
| <b>FC-0</b>            | Lowest layer of Fibre Channel transport. Represents physical media.                                                                                                                                                                    |
| <b>FC-1</b>            | Layer of Fibre Channel transport that contains the 8b/10b encoding scheme.                                                                                                                                                             |
| <b>FC-2</b>            | Layer of Fibre Channel transport that handles framing and protocol, frame format, sequence/exchange management, and ordered set usage.                                                                                                 |

|                   |                                                                                                                                                                                                                                                                                                               |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>FC-3</b>       | Layer of Fibre Channel transport that contains common services used by multiple N_Ports in a node.                                                                                                                                                                                                            |
| <b>FC-4</b>       | Layer of Fibre Channel transport that handles standards and profiles for mapping upper-level protocols such as SCSI and IP onto the Fibre Channel Protocol.                                                                                                                                                   |
| <b>FC-CT</b>      | Fibre Channel common transport.                                                                                                                                                                                                                                                                               |
| <b>FC-FG</b>      | Fibre Channel generic requirements.                                                                                                                                                                                                                                                                           |
| <b>FC-GS</b>      | Fibre Channel generic services.                                                                                                                                                                                                                                                                               |
| <b>FC-GS-2</b>    | Fibre Channel generic services, second generation.                                                                                                                                                                                                                                                            |
| <b>FC-GS-3</b>    | Fibre Channel Generic Services, third generation.                                                                                                                                                                                                                                                             |
| <b>FC_IP</b>      | Fibre Channel-Over-IP.                                                                                                                                                                                                                                                                                        |
| <b>FC-PH</b>      | The Fibre Channel physical and signaling standard for FC-0, FC-1, and FC-2 layers of the Fibre Channel Protocol. Indicates signaling used for cable plants, media types, and transmission speeds.                                                                                                             |
| <b>FCP</b>        | Fibre Channel Protocol. Mapping of protocols onto the Fibre Channel standard protocols. For example, SCSI FCP maps SCSI-3 onto Fibre Channel.                                                                                                                                                                 |
| <b>FCS</b>        | <i>See</i> Fibre Channel Standard.                                                                                                                                                                                                                                                                            |
| <b>FCS switch</b> | Relates to the Brocade Secure Fabric OS feature. One or more designated switches that store and manage security parameters and configuration data for all switches in the fabric. They also act as a set of backup switches to the primary FCS switch. <i>See also</i> <a href="#">“primary FCS switch”</a> . |
| <b>FC-SW-2</b>    | The second-generation Fibre Channel Switch Fabric standard defined by ANSI. Specifies tools and algorithms for the interconnection and initialization of Fibre Channel switches to create a multistage Fibre Channel fabric.                                                                                  |
| <b>FDMI</b>       | Fabric-Device Management Interface. FDMI is a database service provided by the fabric for Nx_Ports. The primary use is by HBA devices that register information about themselves and their ports.                                                                                                             |
| <b>FFFFF5</b>     | Well-known Fibre Channel address for a Class 6 multicast server.                                                                                                                                                                                                                                              |
| <b>FFFFF6</b>     | Well-known Fibre Channel address for a clock synchronization server.                                                                                                                                                                                                                                          |
| <b>FFFFF7</b>     | Well-known Fibre Channel address for a security key distribution server.                                                                                                                                                                                                                                      |
| <b>FFFFF8</b>     | Well-known Fibre Channel address for an alias server.                                                                                                                                                                                                                                                         |
| <b>FFFFF9</b>     | Well-known Fibre Channel address for a QoS facilitator.                                                                                                                                                                                                                                                       |
| <b>FFFFFA</b>     | Well-known Fibre Channel address for a management server.                                                                                                                                                                                                                                                     |
| <b>FFFFFB</b>     | Well-known Fibre Channel address for a time server.                                                                                                                                                                                                                                                           |
| <b>FFFFFC</b>     | Well-known Fibre Channel address for a directory server.                                                                                                                                                                                                                                                      |

|                                |                                                                                                                                                                                                                                                                                                                              |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>FFFFFFD</b>                 | Well-known Fibre Channel address for a fabric controller.                                                                                                                                                                                                                                                                    |
| <b>FFFFFFE</b>                 | Well-known Fibre Channel address for a fabric F_Port.                                                                                                                                                                                                                                                                        |
| <b>FFFFFFF</b>                 | Well-known Fibre Channel address for a broadcast alias ID.                                                                                                                                                                                                                                                                   |
| <b>Fibre Channel</b>           | Fibre Channel is a protocol used to transmit data between servers, switches, and storage devices. It is a high-speed, serial, bidirectional, topology-independent, multiprotocol, and highly scalable interconnection between computers, peripherals, and networks.                                                          |
| <b>FICON®</b>                  | A protocol used on IBM mainframes. Brocade SilkWorm switch FICON® support enables a SilkWorm fabric to transmit FICON® format data between FICON® capable servers and storage.                                                                                                                                               |
| <b>FIFO</b>                    | First in, first out. Refers to a data buffer that follows the first in, first out rule.                                                                                                                                                                                                                                      |
| <b>firmware</b>                | The basic operating system provided with the hardware.                                                                                                                                                                                                                                                                       |
| <b>FL_Port</b>                 | Fabric loop port. A port that is able to transmit under fabric protocol and also has arbitrated loop capabilities. Can be used to connect an NL_Port to a switch. <i>See also</i> “F_Port”, “Fx_Port”.                                                                                                                       |
| <b>flash</b>                   | Programmable nonvolatile RAM (NVRAM) memory that maintains its contents without power.                                                                                                                                                                                                                                       |
| <b>FLOGI</b>                   | Fabric login. The process by which an N_Port determines whether a fabric is present and, if so, exchanges service parameters with it. <i>See also</i> “PLOGI”.                                                                                                                                                               |
| <b>frame</b>                   | The Fibre Channel structure used to transmit data between ports. Consists of a start-of-frame delimiter, header, optional headers, data payload, cyclic redundancy check (CRC), and end-of-frame delimiter. There are two types of frames: link control frames (transmission acknowledgements and so forth) and data frames. |
| <b>FRU</b>                     | Field-replaceable unit. A component that can be replaced onsite.                                                                                                                                                                                                                                                             |
| <b>FSPF</b>                    | Fabric shortest path first. The standard routing protocol for Fibre Channel switches.                                                                                                                                                                                                                                        |
| <b>FSS</b>                     | Fabric OS state synchronization. The FSS service is related to high availability (HA). The primary function of FSS is to deliver state update messages from active components to their peer standby components. FSS determines if fabric elements are synchronized (and thus FSS "compliant").                               |
| <b>FTP</b>                     | File Transfer Protocol.                                                                                                                                                                                                                                                                                                      |
| <b>full fabric</b>             | The Brocade software license that allows multiple E_Ports on a switch, making it possible to create multiple ISL links.                                                                                                                                                                                                      |
| <b>full fabric citizenship</b> | A loop device that has an entry in the Simple Name Server.                                                                                                                                                                                                                                                                   |
| <b>full-duplex</b>             | A mode of communication that allows the same port to simultaneously transmit and receive frames. <i>See also</i> “half-duplex”.                                                                                                                                                                                              |
| <b>Fx_Port</b>                 | A fabric port that can operate as either an F_Port or FL_Port. <i>See also</i> “F_Port”, “FL_Port”.                                                                                                                                                                                                                          |

## G

|                  |                                                                                                                                                                                                     |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>G_Port</b>    | Generic port. A port that can operate as either an E_Port or an F_Port. A port is defined as a G_Port when it is not yet connected or has not yet assumed a specific function in the fabric.        |
| <b>gateway</b>   | Hardware that connects incompatible networks by providing translation for both hardware and software. For example, an ATM gateway can be used to connect a Fibre Channel link to an ATM connection. |
| <b>GBIC</b>      | Gigabit interface converter. A removable serial transceiver module that allows gigabaud physical-level transport for Fibre Channel and gigabit Ethernet.                                            |
| <b>Gbit/sec</b>  | Gigabits per second (1,062,500,000 bits/second).                                                                                                                                                    |
| <b>Gbyte/sec</b> | Gigabytes per second (1,062,500,000 bytes/second).                                                                                                                                                  |
| <b>GLM</b>       | Gigabit Link Module. A semitransparent transceiver that incorporates serializing/deserializing functions.                                                                                           |
| <b>GMT</b>       | Greenwich Mean Time. An international time zone. Also known as "UTC."                                                                                                                               |
| <b>GUI</b>       | A graphic user interface, such as Brocade Web Tools and Brocade Fabric Manager.                                                                                                                     |

## H

|                     |                                                                                                                                                                                                                                     |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>HA</b>           | High availability. A set of features in Brocade SilkWorm switches that is designed to provide maximum reliability and nondisruptive replacement of key hardware and software modules.                                               |
| <b>half-duplex</b>  | A mode of communication that allows a port to either transmit or receive frames at any time except simultaneously (with the exception of link control frames, which can be transmitted at any time). <i>See also "full-duplex".</i> |
| <b>hard address</b> | The AL_PA that an NL_Port attempts to acquire during loop initialization.                                                                                                                                                           |
| <b>HBA</b>          | Host bus adapter. The interface blade between a server or workstation bus and the Fibre Channel network.                                                                                                                            |
| <b>header</b>       | A Fibre Channel frame has a header and a payload. The header contains control and addressing information associated with the frame.                                                                                                 |
| <b>HiPPI</b>        | High-Performance Parallel Interface. An 800 Mb/sec interface normally used in supercomputer environments.                                                                                                                           |
| <b>hop count</b>    | The number of ISLs a frame must traverse to get from its source to its destination.                                                                                                                                                 |
| <b>host</b>         | A computer system that provides end users with services like computation and storage access.                                                                                                                                        |
| <b>HTTP</b>         | Hypertext Transfer Protocol. The standard TCP/IP transfer protocol used on the World Wide Web.                                                                                                                                      |
| <b>hub</b>          | A Fibre Channel wiring concentrator that collapses a loop topology into a physical star topology. Nodes are automatically added to the loop when active and removed when inactive.                                                  |

## I

|                                 |                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>I2C</b>                      | Related to internal circuitry on motherboard. <i>[Is this useful?]</i>                                                                                                                        |
| <b>idle</b>                     | Continuous transmission of an ordered set over a Fibre Channel link when no data is being transmitted, to keep the link active and maintain bit, byte, and word synchronization.              |
| <b>in-band</b>                  | Transmission of management protocol over the Fibre Channel.                                                                                                                                   |
| <b>initiator</b>                | A server or workstation on a Fibre Channel network that initiates communications with storage devices. <i>See also</i> <a href="#">“target”</a> .                                             |
| <b>Insistent Domain ID Mode</b> | Sets the domain ID of a switch as insistent, so that it remains the same over reboots, power cycles, failovers, and fabric reconfigurations. This mode is required to support FICON® traffic. |
| <b>interswitch link</b>         | <i>See</i> <a href="#">“ISL”</a> .                                                                                                                                                            |
| <b>IOCTL</b>                    | I/O control.                                                                                                                                                                                  |
| <b>IP</b>                       | Internet Protocol. The addressing part of TCP.                                                                                                                                                |
| <b>ISL</b>                      | Interswitch link. A Fibre Channel link from the E_Port of one switch to the E_Port of another. <i>See also</i> <a href="#">“E_Port”</a> .                                                     |
| <b>isolated E_Port</b>          | An E_Port that is online but not operational due to overlapping domain IDs or nonidentical parameters (such as E_D_TOVs). <i>See also</i> <a href="#">“E_Port”</a> .                          |
| <b>IU</b>                       | Information unit. A set of information as defined by either an upper-level process protocol definition or upper-level protocol mapping.                                                       |

## J

|             |                                                                                                                                                         |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>JBOD</b> | “Just a bunch of disks.” Indicates a number of disks connected in a single chassis to one or more controllers. <i>See also</i> <a href="#">“RAID”</a> . |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|

## K

|                 |                                                                                                                                                                                                                                                                                  |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>K28.5</b>    | A special 10-bit character used to indicate the beginning of a transmission word that performs Fibre Channel control and signaling functions. The first seven bits of the character are the comma pattern.                                                                       |
| <b>key</b>      | A string of data (usually a numeric value) shared between two entities and used to control a cryptographic algorithm. Usually selected from a large pool of possible keys to make unauthorized identification of the key difficult. <i>See also</i> <a href="#">“key pair”</a> . |
| <b>key pair</b> | In public key cryptography, a pair of keys consisting of an entity's public and private key. The public key can be publicized, but the private key must be kept secret. <i>See also</i> <a href="#">“public key cryptography”</a> .                                              |

## L

|                            |                                                                                                                                                                                                                                                                                                                                              |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>L_Port</b>              | Loop port. A node port (NL_Port) or fabric port (FL_Port) that has arbitrated loop capabilities. An L_Port can be in either Fabric Mode or Loop Mode.                                                                                                                                                                                        |
| <b>LAN</b>                 | Local area network. A network in which transmissions typically take place over fewer than 5 kilometers (3.4 miles).                                                                                                                                                                                                                          |
| <b>latency</b>             | The time required to transmit a frame. Together, latency and bandwidth define the speed and capacity of a link or system.                                                                                                                                                                                                                    |
| <b>LED</b>                 | Light-emitting diode. Used to indicate the status of elements on a switch.                                                                                                                                                                                                                                                                   |
| <b>LIFA</b>                | Loop-initialization fabric-assigned frame. Contains a bitmap of all fabric-assigned AL_PAs and is the first frame transmitted in the loop initialization process after a temporary loop master has been selected.                                                                                                                            |
| <b>LIHA</b>                | Loop-initialization hard-assigned frame. A hard-assigned AL_PA that is indicated by a bit set and is the third frame transmitted in the loop initialization process after a temporary loop master has been selected.                                                                                                                         |
| <b>LILP</b>                | Loop-initialization loop-position frame. The final frame transmitted in a loop initialization process. A returned LIRP contains an accumulation of all of the AL_PA position maps. This allows loop members to determine their relative loop position. This is an optional frame and is not transmitted unless the LIRP is also transmitted. |
| <b>Link Services</b>       | A protocol for link-related actions.                                                                                                                                                                                                                                                                                                         |
| <b>LIP</b>                 | Loop initialization primitive. The signal used to begin initialization in a loop. Indicates either loop failure or node resetting.                                                                                                                                                                                                           |
| <b>LIPA</b>                | Loop-initialization previously assigned. The device marks a bit in the bitmap if it did not log in with the fabric in a previous loop initialization.                                                                                                                                                                                        |
| <b>LIRP</b>                | Loop-initialization report position frame. The first frame transmitted in the loop initialization process after all L_Ports have selected an AL_PA. The LIRP gets transmitted around the loop so all L_Ports can report their relative physical position. This is an optional frame.                                                         |
| <b>LISA</b>                | Loop-initialization soft-assigned frame. The fourth frame transmitted in the loop initialization process after a temporary loop master has been selected. L_Ports that have not selected an AL_PA in a LIFA, LIPA, or LIHA frame select their AL_PA here.                                                                                    |
| <b>LISM</b>                | Loop-initialization select master frame. The first frame transmitted in the initialization process when L_Ports select an AL_PA. LISM is used to select a temporary loop master or the L_Port that will subsequently start transmission of the LIFA, LIPA, LIHA, LISA, LIRP, or LILP frames.                                                 |
| <b>LM_TOV</b>              | Loop master timeout value. The minimum time that the loop master waits for a loop initialization sequence to return.                                                                                                                                                                                                                         |
| <b>loop initialization</b> | The logical procedure used by an L_Port to discover its environment. Can be used to assign AL_PA addresses, detect loop failure, or reset a node.                                                                                                                                                                                            |
| <b>looplest</b>            | A set of devices connected in a loop to a port that is a member of another loop.                                                                                                                                                                                                                                                             |

- LR** Link reset. A primitive sequence used during link initialization between two N\_Ports in point-to-point topology or an N\_Port and an F\_Port in fabric topology. The expected response is an LRR.
- LRR** Link reset response. A primitive sequence during link initialization between two N\_Ports in point-to-point topology or an N\_Port and an F\_Port in fabric topology. It is sent in response to an LR and expects a response of Idle.

## M

- MALLOC** Memory allocation. Usually relates to buffer credits.
- MB/sec** Megabytes per second.
- Mb/sec** Megabits per second.
- metric** A relative value assigned to a route to aid in calculating the shortest path (1000 @ 1 Gbit/sec, 500 @ 2 Gbit/sec).
- MIB** Management Information Base. An SNMP structure to help with device management, providing configuration and device information.
- MRK** Mark primitive signal. Used only in arbitrated loop, MRK is transmitted by an L\_Port for synchronization and is vendor specific.
- MS** Management Server. The Management Server allows a storage area network (SAN) management application to retrieve information and administer the fabric and interconnected elements, such as switches, servers, and storage devices. The MS is located at the Fibre Channel well-known address FFFFFFFAh.
- multicast** The transmission of data from a single source to multiple specified N\_Ports (as opposed to all the ports on the network). *See also* “broadcast”, “unicast”.

## N

- N\_Port** Node port. A port on a node that can connect to a Fibre Channel port or to another N\_Port in a point-to-point connection. *See also* “NL\_Port”, “Nx\_Port”.
- Name Server** Simple Name Server (SNS). A switch service that stores names, addresses, and attributes for up to 15 minutes and provides them as required to other devices in the fabric. SNS is defined by Fibre Channel standards and exists at a well-known address. Also referred to as "directory service."
- NL\_Port** Node loop port. A node port that has arbitrated loop capabilities. Used to connect an equipment port to the fabric in a loop configuration through an FL\_Port. *See also* “N\_Port”, “Nx\_Port”.
- node** A Fibre Channel device that contains an N\_Port or NL\_Port.
- node count** The number of nodes attached to a fabric.
- node name** The unique identifier for a node, communicated during login and port discovery.

**NOS** Not operational. The NOS primitive sequence is transmitted to indicate that the FC\_Port transmitting the NOS has detected a link failure or is offline, waiting for the offline sequence (OLS) to be received.

**NS** Name Server. The service provided by a fabric switch that stores names, addresses, and attributes related to Fibre Channel objects. Can cache information for up to 15 minutes. Also known as “Simple Name Server” or as a “directory service.” *See also* “[SNS](#)”.

**Nx\_Port** A node port that can operate as either an N\_Port or NL\_Port.

## O

**OLS** Primitive sequence offline.

**ON** Offline notification. Refers to an ELS field that appears in **portlogdump** command output.

**OPN** Open primitive signal. Applies only to arbitrated loop; sent by an L\_Port that has won the arbitration process to open communication with one or more ports on the loop.

**ordered set** A transmission word that uses 8B/10B mapping and begins with the K28.5 character. Ordered sets occur outside of frames and include the following items:

**Frame delimiters.** Mark frame boundaries and describe frame contents.

**Primitive signals.** Indicate events.

**Primitive sequences.** Indicate or initiate port states.

Ordered sets are used to differentiate Fibre Channel control information from data frames and to manage frame transport.

**originator** The Nx\_Port that originated an exchange.

**out-of-band** Transmission of management protocol outside of the Fibre Channel network, usually over Ethernet.

**OX\_ID** Originator ID. Refers to the exchange ID assigned by the originator port.

## P

**parallel** The simultaneous transmission of data bits over multiple lines.

**path selection** The selection of a transmission path through the fabric. Brocade switches use the FSPF protocol. *See also* “[FSPF](#)”.

**payload** A Fibre Channel frame has a header and a payload. The payload contains the information being transported by the frame; it is determined by the higher-level service or FC\_4 upper-level protocol. There are many different payload formats, based on protocol and size of truck bed.

**Performance Monitoring** A Brocade SilkWorm switch feature that monitors port traffic and includes frame counters, SCSI read monitors, SCSI write monitors, and other types of monitors.

**phantom address** An AL\_PA value that is assigned to a device that is not physically in the loop. Also known as “phantom AL\_PA.”

|                                  |                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>phantom device</b>            | A device that is not physically in an arbitrated loop but is logically included through the use of a phantom address.                                                                                                                                                                                                                           |
| <b>PID</b>                       | Port identifier.                                                                                                                                                                                                                                                                                                                                |
| <b>PKI</b>                       | Public key infrastructure. An infrastructure that is based on public key cryptography and CA (certificate authority) and that uses digital certificates. <i>See also</i> “CA”, “digital certificate”, “public key cryptography”.                                                                                                                |
| <b>PKI certification utility</b> | Public key infrastructure certification utility. A utility that makes it possible to collect certificate requests from switches and to load certificates to switches. <i>See also</i> “digital certificate”, “PKI”.                                                                                                                             |
| <b>PLOGI</b>                     | Port login. The port-to-port login process by which initiators establish sessions with targets. <i>See also</i> “FLOGI”.                                                                                                                                                                                                                        |
| <b>point-to-point</b>            | A Fibre Channel topology that employs direct links between each pair of communicating entities. <i>See also</i> “topology”.                                                                                                                                                                                                                     |
| <b>port</b>                      | In a Brocade SilkWorm switch environment, an SFP or GBIC receptacle on a switch to which an optic cable for another device is attached.                                                                                                                                                                                                         |
| <b>port address</b>              | In Fibre Channel technology, the port address is defined in hexadecimal. In the Brocade Fabric OS, a port address can be defined by a domain and port number combination or by area number. In an ESCON Director, an address used to specify port connectivity parameters and to assign link addresses for attached channels and control units. |
| <b>port blade</b>                | A hardware component that provides a platform for field-replaceable, hot-swappable ports.                                                                                                                                                                                                                                                       |
| <b>port log</b>                  | A record of all activity on a switch, kept in volatile memory.                                                                                                                                                                                                                                                                                  |
| <b>port log dump</b>             | A view of what happens on a switch, from the switch's point of view. The <b>portlogdump</b> command is used to read the port log.                                                                                                                                                                                                               |
| <b>port name</b>                 | A user-defined alphanumeric name for a port.                                                                                                                                                                                                                                                                                                    |
| <b>port swapping</b>             | Port swapping is the ability to redirect a failed port to another port. This feature is available in Fabric OS v4.1.0 and higher.                                                                                                                                                                                                               |
| <b>port_name</b>                 | The unique identifier assigned to a Fibre Channel port. Communicated during login and port discovery.                                                                                                                                                                                                                                           |
| <b>POST</b>                      | Power-on self-test. A series of tests run by a switch after it is turned on.                                                                                                                                                                                                                                                                    |
| <b>primary FCS switch</b>        | Relates to the Brocade Secure Fabric OS feature. The primary fabric configuration server switch actively manages security and configurations for all switches in the fabric.                                                                                                                                                                    |
| <b>primitive sequence</b>        | An ordered set that is transmitted repeatedly and continuously. Primitive sequences are transmitted to indicate specific conditions within or conditions encountered by the receiver logic of an FC_Port. <i>See</i> “OLS” and “NOS”.                                                                                                           |
| <b>primitive signals</b>         | An ordered set that indicates actions or events and requires just one occurrence to trigger a response. Idle and R_RDY are used in all three topologies: ARB, OPN, and CLS. MRK is used in arbitrated loop.                                                                                                                                     |

|                                |                                                                                                                                                                                                                                                                                                 |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>principal switch</b>        | The first switch to boot up in a fabric. Ensures unique domain IDs among roles.                                                                                                                                                                                                                 |
| <b>private key</b>             | The secret half of a key pair. <i>See also</i> “key”, “key pair”.                                                                                                                                                                                                                               |
| <b>private loop</b>            | An arbitrated loop that does not include a participating FL_Port.                                                                                                                                                                                                                               |
| <b>private loop device</b>     | A device that supports a loop and can understand 8-bit addresses but does not log in to the fabric.                                                                                                                                                                                             |
| <b>private NL_Port</b>         | An NL_Port that communicates only with other private NL_Ports in the same loop and does not log in to the fabric.                                                                                                                                                                               |
| <b>protocol</b>                | A defined method and set of standards for communication. Determines the type of error-checking, the data-compression method, how sending devices indicate an end of message, and how receiving devices indicate receipt of a message.                                                           |
| <b>pstate</b>                  | Port State Machine.                                                                                                                                                                                                                                                                             |
| <b>public device</b>           | A device that supports arbitrated loop protocol, can interpret 8-bit addresses, and can log in to the fabric.                                                                                                                                                                                   |
| <b>public key</b>              | The public half of a key pair. <i>See also</i> “key”, “key pair”.                                                                                                                                                                                                                               |
| <b>public key cryptography</b> | A type of cryptography that uses a key pair, with the two keys in the pair called at different points in the algorithm. The sender uses the recipient's public key to encrypt the message, and the recipient uses the recipient's private key to decrypt it. <i>See also</i> “key pair”, “PKI”. |
| <b>public loop</b>             | An arbitrated loop that includes a participating FL_Port and can contain both public and private NL_Ports.                                                                                                                                                                                      |
| <b>public NL_Port</b>          | An NL_Port that logs in to the fabric, can function within either a public or a private loop, and can communicate with either private or public NL_Ports.                                                                                                                                       |

## Q

|              |                                                                                                         |
|--------------|---------------------------------------------------------------------------------------------------------|
| <b>QoS</b>   | Quality of service.                                                                                     |
| <b>quad</b>  | A group of four adjacent ports that share a common pool of frame buffers.                               |
| <b>queue</b> | A mechanism for each AL_PA address that allows for collecting frames prior to sending them to the loop. |

## R

|                |                                                                                                                                                       |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>R_A_TOV</b> | Resource allocation timeout value. The maximum time a frame can be delayed in the fabric and still be delivered. <i>See also</i> “E_D_TOV”, “RR_TOV”. |
| <b>R_CTL</b>   | Route control. The first 8 bits of the header, which defines the type of frame and its contents.                                                      |

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>R_RDY</b>             | Receiver ready. A primitive signal indicating that the port is ready to receive a frame.                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>R_T_TOV</b>           | Receiver transmitter timeout value, used by receiver logic to detect loss of synchronization between transmitters and receivers.                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>RAID</b>              | Redundant array of independent disks. A collection of disk drives that appear as a single volume to the server and are fault tolerant through mirroring or parity checking. <i>See also</i> “JBOD”.                                                                                                                                                                                                                                                                                                                                          |
| <b>RCS</b>               | Reliable Commit Service. Refers to Brocade-specific ILS command code.                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>remote switch</b>     | An optional product for long-distance fabrics, requiring a Fibre Channel-to-ATM or SONET gateway.                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>responder</b>         | The N_Port with which an exchange originator wishes to communicate.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>RLS</b>               | Read Link Status.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>route</b>             | As it applies to a fabric, the communication path between two switches. Might also apply to the specific path taken by an individual frame, from source to destination. <i>See also</i> “FSPF”.                                                                                                                                                                                                                                                                                                                                              |
| <b>routing</b>           | The assignment of frames to specific switch ports, according to frame destination.                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>RR_TOV</b>            | Resource recovery timeout value. The minimum time a target device in a loop waits after a LIP before logging out an SCSI initiator. <i>See also</i> “E_D_TOV”, “R_A_TOV”.                                                                                                                                                                                                                                                                                                                                                                    |
| <b>RSCN</b>              | Registered state change notification. A switch function that allows notification of fabric changes to be sent from the switch to specified nodes. The fabric controller issues RSCN requests to N_Ports and NL_Ports, but only if they have registered to be notified of state changes in other N_Ports and NL_Ports. This registration is performed via the State Change Registration (SCR) Extended Link Service. An N_Port or NL_Port can issue an RSCN to the fabric controller without having completed SCR with the fabric controller. |
| <b>RTWR</b>              | Reliable transport with response. Might appear as a task in <b>portlogdump</b> command output.                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>running disparity</b> | A binary parameter indicating the cumulative disparity (positive or negative) of all previously issued transmission characters.                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>RW</b>                | Read/write. Refers to access rights.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>RX</b>                | Receiving frames.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>RX_ID</b>             | Responder exchange identifier. A 2-byte field in the frame header that can be used by the responder of the exchange to identify frames as being part of a particular exchange.                                                                                                                                                                                                                                                                                                                                                               |

## S

|             |                                                                                                                                          |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------|
| <b>S_ID</b> | Source ID. Refers to the native port address (24 bit address).                                                                           |
| <b>SAN</b>  | Storage area network. A network of systems and storage devices that communicate using Fibre Channel protocols. <i>See also</i> “fabric”. |

|                                 |                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>SAN architecture</b>         | The overall design of a storage network solution, which includes one or more related fabrics, each of which has a topology.                                                                                                                                                              |
| <b>SAN port count</b>           | The number of ports available for connection by nodes in the entire SAN.                                                                                                                                                                                                                 |
| <b>SCN</b>                      | State change notification. Used for internal state change notifications, not external changes. This is the switch logging that the port is online or is an Fx_port, not what is sent from the switch to the Nx_ports.                                                                    |
| <b>SCR</b>                      | State change registration. Extended Link Service (ELS) requests the fabric controller to add the N_Port or NL_Port to the list of N_Ports and NL_Ports registered to receive the Registered State Change Notification (RSCN) Extended Link Service.                                      |
| <b>SCSI</b>                     | Small Computer Systems Interface. A parallel bus architecture and a protocol for transmitting large data blocks to a distance of 15 to 25 meters.                                                                                                                                        |
| <b>sectelnet</b>                | A protocol similar to telnet but with encrypted passwords for increased security.                                                                                                                                                                                                        |
| <b>Secure Fabric OS</b>         | A separately sold Brocade feature that provides advanced, centralized security for a fabric.                                                                                                                                                                                             |
| <b>security policy</b>          | Rules that determine how security is implemented in a fabric. Security policies can be customized through Brocade Secure Fabric OS or Brocade Fabric Manager.                                                                                                                            |
| <b>SEQ_ID</b>                   | Sequence identifier. A 1-byte field in the frame header change to identify the frames as being part of a particular exchange sequence between a pair of ports.                                                                                                                           |
| <b>sequence</b>                 | A group of related frames transmitted in the same direction between two N_Ports.                                                                                                                                                                                                         |
| <b>sequence initiator</b>       | The N_Port that begins a new sequence and transmits frames to another N_Port.                                                                                                                                                                                                            |
| <b>sequence recipient</b>       | Serializing/deserializing circuitry. A circuit that converts a serial bit stream into parallel characters, and vice-versa.                                                                                                                                                               |
| <b>SES</b>                      | SCSI Enclosure Services. A subset of the SCSI protocol used to monitor temperature, power, and fan status for enclosed devices.                                                                                                                                                          |
| <b>SFP</b>                      | Small-form-factor pluggable. A transceiver used on 2 Gbit/sec or 4Gbit/sec switches that replaces the GBIC.                                                                                                                                                                              |
| <b>SilkWorm</b>                 | The brand name for the Brocade family of switches.                                                                                                                                                                                                                                       |
| <b>Simple Name Server (SNS)</b> | A switch service that stores names, addresses, and attributes for up to 15 minutes and provides them as required to other devices in the fabric. SNS is defined by Fibre Channel standards and exists at a well-known address. Also referred to as "directory service" or "name server." |
| <b>Single CP Mode</b>           | The <b>-s</b> option of the <b>firmwaredownload</b> command. Using <b>firmwaredownload -s</b> enables Single CP Mode. In the SilkWorm 12000, Single CP Mode enables a user to upgrade a single CP blade and to select full-install, auto-reboot, and auto-commit.                        |

|                                        |                                                                                                                                                                                                                                                                                                                |
|----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>SNMP</b>                            | Simple Network Management Protocol. An Internet management protocol that uses either IP for network-level functions and UDP for transport-level functions, or TCP/IP for both. Can be made available over other protocols, such as UDP/IP, because it does not rely on the underlying communication protocols. |
| <b>SNS</b>                             | Simple Name Server.                                                                                                                                                                                                                                                                                            |
| <b>SOF</b>                             | Start of frame. A group of ordered sets that marks the beginning of a frame and indicates the class of service the frame will use.                                                                                                                                                                             |
| <b>SONET</b>                           | Synchronous optical network. A standard for optical networks that provides building blocks and flexible payload mappings.                                                                                                                                                                                      |
| <b>special character</b>               | A 10-bit character that does not have a corresponding 8-bit value but is still considered valid. The special character is used to indicate that a particular transmission word is an ordered set. This is the only type of character to have five 1s or 0s in a row.                                           |
| <b>SSH</b>                             | Secure shell. Used starting in Brocade Fabric OS v4.1.0 to support encrypted telnet sessions to the switch. SSH encrypts all messages, including the client sending the password at login.                                                                                                                     |
| <b>switch</b>                          | A fabric device providing bandwidth and high-speed routing of data via link-level addressing.                                                                                                                                                                                                                  |
| <b>switch name</b>                     | The arbitrary name assigned to a switch.                                                                                                                                                                                                                                                                       |
| <b>switch port</b>                     | A port on a switch. Switch ports can be E_Ports, F_Ports, or FL_Ports.                                                                                                                                                                                                                                         |
| <b>switch-to-switch authentication</b> | The process of authenticating both switches in a switch-to-switch connection using digital certificates. <i>See also</i> <a href="#">“authentication”</a> , <a href="#">“digital certificate”</a> .                                                                                                            |
| <b>syslog</b>                          | Syslog daemon. Used to forward error messages.                                                                                                                                                                                                                                                                 |

## T

|                    |                                                                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>T11</b>         | A standards committee chartered with creating standards for Fibre Channel.                                                                                                    |
| <b>target</b>      | A storage device on a Fibre Channel network. <i>See also</i> <a href="#">“initiator”</a> .                                                                                    |
| <b>TC</b>          | Track changes.                                                                                                                                                                |
| <b>telnet</b>      | A virtual terminal emulation used with TCP/IP. "Telnet" is sometimes used as a synonym for the Brocade Fabric OS CLI.                                                         |
| <b>tenancy</b>     | The time from when a port wins arbitration in a loop until the same port returns to the monitoring state. Also referred to as "loop tenancy."                                 |
| <b>Time Server</b> | A Fibre Channel service that allows for the management of all timers.                                                                                                         |
| <b>topology</b>    | As it applies to Fibre Channel technology, the configuration of the Fibre Channel network and the resulting communication paths allowed. There are three possible topologies: |

**Point to point.** A direct link between two communication ports.

**Switched fabric.** Multiple N\_Ports linked to a switch by F\_Ports.

**Arbitrated loop.** Multiple NL\_Ports connected in a loop.

|                               |                                                                                                                                                                                                                                             |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>track changes</b>          | A Brocade Fabric OS feature that can be enabled to report specific activities (for example, logins, logouts, and configuration task changes). The output from the track-changes feature is dumped to the system message log for the switch. |
| <b>transceiver</b>            | A device that converts one form of signaling to another for transmission and reception; in fiber optics, optical to electrical.                                                                                                             |
| <b>Translative Mode</b>       | A mode in which private devices can communicate with public devices across the fabric.                                                                                                                                                      |
| <b>transmission character</b> | A 10-bit character encoded according to the rules of the 8B/10B algorithm.                                                                                                                                                                  |
| <b>transmission word</b>      | A group of four transmission characters.                                                                                                                                                                                                    |
| <b>trap (SNMP)</b>            | The message sent by an SNMP agent to inform the SNMP management station of a critical error. <i>See also</i> “SNMP”.                                                                                                                        |
| <b>trunking</b>               | In Fibre Channel technology, a feature that enables distribution of traffic over the combined bandwidth of up to four ISLs between adjacent switches, while preserving in-order delivery.                                                   |
| <b>trunking group</b>         | A set of up to four trunked ISLs.                                                                                                                                                                                                           |
| <b>trunking ports</b>         | The ports in a set of trunked ISLs.                                                                                                                                                                                                         |
| <b>TS</b>                     | Time Server.                                                                                                                                                                                                                                |
| <b>TTL</b>                    | Time-to-live. The number of seconds an entry exists in cache before it expires.                                                                                                                                                             |
| <b>tunneling</b>              | A technique for enabling two networks to communicate when the source and destination hosts are both on the same type of network but are connected by a different type of network.                                                           |
| <b>TX</b>                     | Transmit.                                                                                                                                                                                                                                   |

## U

|                |                                                                                                                                                                                                                                                        |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>U_Port</b>  | Universal port. A switch port that can operate as a G_Port, E_Port, F_Port, or FL_Port. A port is defined as a U_Port when it is not connected or has not yet assumed a specific function in the fabric. <i>[How is this different from a G_Port?]</i> |
| <b>unicast</b> | The transmission of data from a single source to a single destination. <i>See also</i> “broadcast”, “multicast”.                                                                                                                                       |
| <b>UTC</b>     | Universal Time Conversion. Also known as “Coordinated Universal Time,” which is an international standard of time. UTC is 8 hours behind Pacific Standard Time and 5 hours behind Eastern Standard Time. <i>See also</i> “GMT”.                        |

## W

|                           |                                                                                                                                                               |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>WAN</b>                | Wide area network.                                                                                                                                            |
| <b>watchdog</b>           | A software daemon that monitors Fabric OS modules on the kernel.                                                                                              |
| <b>well-known address</b> | As it pertains to Fibre Channel technology, a logical address defined by Fibre Channel standards as assigned to a specific function and stored on the switch. |
| <b>WWN</b>                | World Wide Name. An identifier that is unique worldwide. Each entity in a fabric has a separate WWN.                                                          |

## Z

|                           |                                                                                                                                                                                                                        |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>zone</b>               | A set of devices and hosts attached to the same fabric and configured as being in the same zone. Devices and hosts within the same zone have access to others in the zone but are not visible to any outside the zone. |
| <b>zone configuration</b> | A specified set of zones. Enabling a configuration enables all zones in that configuration. <i>See also</i> <a href="#">“defined zone configuration”</a> .                                                             |
| <b>zoning</b>             | A feature in fabric switches or hubs that allows segmentation of a node by physical port, name, or address.                                                                                                            |



# Index

---

## A

- accessing switches and fabrics 3-5
- account ID 2-1
- account privilege levels 2-3
- activating
  - a switch certificate 3-22
  - ports on demand 2-13
- adding
  - a new switch or fabric 13-25
  - alias members 13-12
  - and removing FICON CUP licenses 7-14
  - custom filter-based monitors 14-8
  - end-to-end monitors 14-4
  - filter-based monitors 14-8
  - members to a zone configuration 13-19
  - standard filter-based monitors 14-8
  - switches to a zone 13-25
  - zone members 13-14
- advanced performance monitoring commands 14-2
- AIX procedure, PID A-18
- alias
  - adding members 13-12
  - creating 13-12
  - deleting 13-13
  - removing members 13-12
- analyzing connection problems 10-4
- assigning a static route 6-3

## B

- backing up
  - a configuration 4-1
  - and restoring configurations, FICON 7-15
- basic
  - card management 5-3
  - PID procedure A-15
- beaconing mode 5-12
- blocking listeners 3-5

- boot PROMpassword 3-34

- Brocade

- switch requirements for interoperability B-2
  - Web site 2-8

- browser and Java support 3-18

- buffer-limited port 11-1

## C

- card management 5-3
- changes to configuration data A-3
- changing
  - an account password 3-8
  - to core PID format A-9
  - to extended edge PID format A-10
- chassisshow command 2-19
- checking
  - connected switches 4-6
  - status 2-18
- choosing
  - a CA 3-20
  - an extended ISL mode 11-3
- clearing
  - FICON management database 7-3, 7-8
  - monitor counters 14-14
- CLI 1-2
- collecting performance data 14-15
- combining SilkWorm 12000 and 24000 cards in one chassis 5-12

- command
    - advanced performance monitoring 14-2
    - chassisshow 2-19
    - configupload 4-5
    - fabricshow 2-19
    - hashow 2-19
    - licenseadd 2-9
    - licenseremove 2-10
    - licenseshow 2-10
    - nsallshow 2-19
    - nsshow 2-19
    - slotshow 2-19
    - switchshow 2-18
  - configuration
    - FICON environment switched point-to-point 7-2
    - FICON environment, cascaded 7-2
    - high-integrity fabric 7-5
    - new SilkWorm 24000 with two domains 5-9
    - recommendations for interoperability B-3
    - restrictions for interoperability B-3
    - save to a host 4-1
    - settings, FICON environment 7-4
  - configuring a single switch 7-5
  - configuring an extended ISL 11-7
  - configuring for SNMP 3-25
  - configuring for syslogd 9-12
  - configuring for the SSL protocol 3-18
  - configuring secure file copy 3-34
  - configuring switches 7-4
  - configuring the browser 3-23
  - configuring the host 9-13
  - configuring the radius server 3-10
  - configuring the server database 8-5
  - configuring the switch 3-15
  - configuring the switch in a FICON environment 7-4
  - configuring the telnet interface 3-3
  - connecting to devices 2-15
  - connecting to other switches 2-15
  - connecting to the command line interface 2-1
  - connection
    - serial 2-1
    - telnet 2-1
  - conserving power 5-4
  - controlling access 8-2
  - controlling topology discovery 8-6
  - converting an installed SilkWorm 24000 to support two domains 5-10
  - core/edge topology and ISL trunking 12-3
  - correcting device login issues 10-14
  - correcting I2C bus errors 10-12
  - correcting link failures 10-19
  - correcting marginal links 10-22
  - correcting zoning setup issues 10-10
  - CRC errors, displaying 14-3
  - creating a zone 13-14
  - creating a zone configuration 13-18
  - creating an alias 13-12
  - creating and maintaining user-defined accounts 3-6
  - creating and maintaining zones 13-14
  - creating and managing zone aliases 13-11
  - creating and modifying zoning configurations 13-18
  - customizing the chassis name 2-12
  - customizing the switch name 2-10
- ## D
- database, clearing in a FICON environment 7-3
  - date and time 2-5
  - default names 2-10
  - default password 2-4
  - defined zone configuration 13-6
  - deleting
    - end-to-end monitors 14-7
    - filter-based monitors 14-10
  - deleting a zone 13-15
  - deleting a zone configuration 13-20
  - deleting an alias 13-13
  - deleting end-to-end monitors 14-7
  - deleting filter-based monitors 14-10
  - designing fabric for trunking 12-3
  - deskew values
    - displaying 12-8
  - device
    - connecting 2-15
  - device-based routing 6-1

- disabled zone configuration 13-6
- disabling and enabling a port 2-12
- disabling and enabling a switch 2-12
- disabling and enabling cards 5-3
- disabling interoperability mode B-6
- displaying
  - CRC error count 14-3
  - end-to-end mask 14-6
  - node identification data, FICON environments 7-7
  - registered listeners for link incidents, FICON environment 7-7
- displaying additional help topics 1-4
- displaying and clearing the CRC error count 14-3
- displaying and deleting certificates 3-24
- displaying command help 1-3
- displaying configuration settings 4-1
- displaying information 7-7
- displaying mode register bit settings 7-12
- displaying monitor counters 14-11
- displaying the fmsmode setting 7-11
- displaying trunking information 12-8
- domain ID
  - FICON environment 2-15
- domain ID, insistent 7-1
- downloading configurations across a fabric 4-5
- downloading firmware 4-7

## E

- effective zone configuration 13-6
- enabling and disabling FICON management
  - server mode 7-9
- enabling and disabling interoperability mode B-5
- enabling and disabling ISL trunking 12-6
- enabling and disabling local authentication 3-18
- enabling and disabling the platform services 8-1
- enabling interoperability mode B-5
- encryption 3-19
- end-to-end monitoring 14-3

- end-to-end monitors
  - adding 14-4
  - deleting 14-7
  - displaying the mask 14-6
  - restoring configuration 14-15
  - saving configuration 14-15
  - setting a mask 14-6
- ensuring network security 3-3
- evaluating the fabric A-5
- example
  - chassisshow 2-19
  - fabricshow 2-19
  - nsallshow 2-19
  - slotshow 2-19
- exchange-based routing 6-2
- extended link buffer allocation 11-1

## F

- fabric
  - high integrity 7-2
- fabric access 3-5
- fabric connectivity 2-19
- fabric considerations 11-2, 12-3
- fabric, designing for trunking 12-3
- fabricshow command 2-19
- fans, status of 9-8
- feature licenses 2-8
- FICON environment
  - cascaded configuration 7-2
  - changing domain id 2-15
  - configuration settings 7-4
  - configuring switches 7-4
  - disabling IDID mode 7-3
  - displaying
    - link incidents 7-3
    - registered listeners for link incidents 7-7
  - enabling IDID mode 7-3
  - high integrity fabric 7-2
  - identifying port swapping nodes 7-8
  - monitoring FRU failures 7-7
  - node identification data, displaying 7-7
  - switched point-to-point configuration 7-2

- filter-based monitors 14-7
  - adding 14-8
  - deleting 14-10
  - restoring configuration 14-15
  - saving configuration 14-15
- firmware download 4-7
- frame transfer with brocade remote switch C-1
- fru failures 7-7
- fru failures, monitoring in FICON environments 7-3, 7-7

## G

- gateway C-1
- gateway, remote switch C-1
- gathering information for technical support 10-2
- generating
  - batch of licenses 2-8
- generating a public/private key 3-20
- generating and storing a csr 3-20

## H

- hard zoning 13-7
- hardware-enforced zoning 13-7
- hashow command 2-19
- help information 1-3
- high availability (HA) 2-19
- high integrity fabric 7-2
- host reboots A-2
- host-based zoning 13-3
- HP/UX procedure A-16
- hybrid update A-9

## I

- ID, account 2-1
- identifying
  - ports from the tag field 7-16
- identifying media-related issues 10-16
- Identifying ports
  - by slot and port number 5-2

- identifying ports 5-1
  - by port area ID 5-2
- IDID 7-1
- IDID mode
  - enabling and disabling in a FICON environment 7-6
- impact of changing the fabric PID format A-2
- inaccurate information in the system message log 10-23
- initializing trunking on ports 12-4
- insistent domain ID 7-1
- installing a root certificate to the Java plug-in 3-24
- installing a switch certificate 3-22
- intermix mode 7-1
- interoperability, Brocade switch requirements B-2
- interswitch link 2-15
- ISL 2-15

## J

- Java support, SSL 3-18

## K

- key
  - license, generating on the web 2-8
  - transaction, for licensed features 2-8

## L

- license key
  - activating 2-9
- license keys
  - generating 2-8
- licenseadd command 2-9
- licensed features 2-8
- licenseremove command 2-10
- licenseshow command 2-9
- link incidents
  - displaying in a FICON environment 7-3, 7-7

- linking through a gateway 2-17
- listing link characteristics 12-9
- login
  - switch 2-1
- long distance ISLs 11-2
- LUN masking 13-3
- LWL
  - ISL Trunking support for 12-1

## M

- maintaining configurations 4-1
- maintaining firmware 4-5
- maintaining licensed features 2-8
- making basic connections 2-14
- managing zoning configurations in a fabric 13-22
- mask for end-to-end monitors
  - displaying 14-6
  - setting 14-6
- monitoring end-to-end performance 14-3
- monitoring filter-based performance 14-7
- monitoring ISL performance 14-10
- monitoring traffic 12-4
- monitoring trunks 14-10
- monitors
  - clearing counters 14-14
- most common problem areas 10-1

## N

- name server zoning 13-3
- network security 3-3
- node identification data 7-7
- nsallshow command 2-19
- nsshow command 2-19

## O

- obtaining and unzipping firmware 4-6
- obtaining certificates 3-21
- obtaining slot information 5-8

- offline update A-8
- online update A-8
- optimizing resources through zoning 13-1

## P

- password 2-1
  - boot prom 3-34
  - default 2-4
- password management information D-1
- password migration during firmware changes D-4
- password prompting behaviors D-3
- password recovery options D-5
- passwords
  - recovering forgotten passwords 3-39
- perfaddeemonitor command 14-4
- perfaddIPmonitor command 14-8
- perfaddusermonitor command 14-8
- perfcfgrestore command 14-15
- perfcfgsave command 14-15
- perfdleemonitor command 14-7
- perfdelfiltermonitor command 14-10
- performance monitoring commands 14-2
- performing PID format changes A-12
- perfsetporteemask command 14-6
- perfshowalpacrc command 14-3
- perfshowporteemask command 14-6
- persistently enabling/disabling ports 7-13
- PID
  - AIX procedure A-18
  - binding A-1
- PKI 3-18
- planning the update procedure A-7
- policies, routing 6-1
- port
  - buffer-limited 11-1
- port and switch naming standards 7-14
- port swapping nodes, identifying in FICON environments 7-8
- port-based routing 6-2

- ports
  - identifying by port area ID 5-2
  - identifying by slot and port number 5-2
  - status of 9-5
- ports, swapping 7-8
- powering off a card 5-3
- powering port cards on and off 5-3
- preparing a switch 7-4
- printing hard copies of switch information 4-5
- privileges in accounts 2-3
- procedural differences between fixed-port and variable-port switches 1-1
- public key infrastructure encryption 3-18

## R

- recognizing buffer underallocation 12-10
- recognizing MQ-WRITE errors 10-12
- recognizing the port initialization and FCP auto discovery process 10-24
- recording configuration information 7-17
- recovering forgotten passwords 3-39
- recovery password 3-37
- recovery string, boot PROM password 3-34
- registered listeners 7-7
- remote switch C-1
- removing
  - end-to-end monitors 14-7
  - filter-based monitors 14-10
- removing members
  - zone 13-15
- removing members from a zone configuration 13-19
- removing members, alias 13-12
- resolving zone conflicts 13-27
- restoring a configuration 4-2
- restoring a segmented fabric 10-8
- restoring configurations in a FICON environment 4-4
- restoring monitor configuration 14-15
- restoring the system configuration settings 4-2
- routing
  - assigning a static route 6-3
  - routing policies 6-1, 6-2
  - rules for configuring zones 13-10

## S

- saved zone configuration 13-6
- saving and restoring monitor configurations 14-15
- saving monitor configuration 14-15
- scope and references 1-2
- secure shell (ssh) 3-3
- secure socket layer protocol 3-18
- secure sockets layer 3-18
- security 3-3
- security and zoning 13-27
- selecting a PID format A-4
- serial connection 2-1, 2-2
- setting a mask for end-to-end monitors 14-6
- setting a unique domain id 7-6
- setting chassis configurations 5-5
- setting mask for end-to-end monitors 14-6
- setting mode register bits 7-13
- setting port speeds 12-7
- setting the boot prom password 3-34
- setting the card beacon mode 5-12
- setting the date and time 2-5
- setting the default account passwords 2-3
- setting the IP address 2-3
- setting the security level 3-26
- setting the switch date and time 2-5
- setting up automatic trace dump transfers 9-15
- setting up RADIUS AAA service 3-9
- setup summary 7-9
- slotshow command 2-19
- SNMP 3-25, 3-26, 3-30, 3-31, 3-32
- SNMPv1 3-28, 3-30
- SNMPv3 3-27
- specifying frame order delivery 6-3
- specifying the routing policy 6-2
- splitting a fabric 13-27
- SSL 3-18
- standard filter-based monitors 14-8

- standard trunking criteria 12-2
- static PID mapping errors A-3
- static route 6-3
- storage-based zoning 13-3
- summary of PID formats A-1
- summary of SSL procedures 3-19
- supported brocade features B-2
- supportsave command 9-14
- swapping port area IDs A-19
- swapping ports 7-8
- switch
  - configuring in a FICON environment 7-4
  - system status 9-3
- switch access 3-5
- switch names 2-10
- switchshow command 2-18
- SWL, ISL Trunking support for 12-1

## T

- tag field, interpreting 7-16
- telnet connection 2-1
- temperature, status of 9-9
- time and date 2-5
- tracking and controlling switch changes 2-20
- traffic patterns
  - planning for 12-3
- transaction key 2-8
- troubleshooting 7-15
- troubleshooting certificates 3-25
- troubleshooting firmware downloads 4-15
- troubleshooting trunking problems 12-9
- trunking
  - displaying information 12-8
  - over distance 11-8
- trunking over extended fabrics 12-9

## U

- unsupported brocade features B-2
- upgrading SilkWorm 12000 and 24000 directors 4-11

- upgrading SilkWorm 3016, 3250, 3850, 3900, and 4100 switches 4-8
- user-defined filter-based monitors 14-8
- using dynamic load sharing 6-4
- using FICON CUP 7-9
- using legacy commands for SNMPv1 3-30
- using the snmpconfig command 3-27
- using zoning to administer security 13-27

## V

- vendor switch requirements B-1
- verify
  - device connectivity 2-15
  - fabric connectivity 2-19
  - high availability (HA) 2-19
- viewing
  - fan status 9-8
  - port status 9-5
  - power supply status 9-9
  - temperature status 9-9
- viewing an alias
  - alias
    - viewing 13-13
- viewing and saving diagnostic information 9-14
- viewing equipment status 9-8
- viewing port information 9-4
- viewing power-on self test 9-1
- viewing routing information along a path 6-7
- viewing routing path information 6-5
- viewing switch status 9-3
- viewing the port log 9-10
- viewing the system message log 9-10
- viewing zone database configurations 13-22
- viewing zones 13-15

## W

- with a recovery string 3-35
- without a recovery string 3-37
- working with domain IDs 2-15

## Z

### zone

- adding members 13-14
- adding switches 13-25
- creating 13-14
- creating a configuration 13-18
- deleting 13-15
- deleting a configuration 13-20
- removing members 13-15
- viewing 13-15
- viewing configurations 13-22

### zone aliases 13-5

### zone configuration

- adding members 13-19
- removing members 13-19

### zone configurations 13-5

### zone name restrictions B-4

### zone objects 13-5

### zone types 13-2

### zoning

- administering security 13-27

### zoning and PDCM considerations 7-15

### zoning concepts 13-2

### zoning enforcement 13-6

### zoning restrictions B-4

### zoning schemes 13-6

### zoning terminology 13-1